

(19)



Europäisches Patentamt

European Patent Office

Office européen des brevets



(11)

EP 0 411 805 B1

(12)

## EUROPEAN PATENT SPECIFICATION

(45) Date of publication and mention  
of the grant of the patent:

09.10.1996 Bulletin 1996/41

(51) Int Cl. 6: G06F 11/16

(21) Application number: 90308002.6

(22) Date of filing: 20.07.1990

### (54) Bulk memory transfer during resync

Massenspeicherübertragung während der Neusynchronisierung

Transfert de mémoire de masse pendant la resynchronisation

(84) Designated Contracting States:

AT BE CH DE DK ES FR GB GR IT LI LU NL SE

• Thirumalai, Ajai

Marlboro, Massachusetts 01752 (US)

(30) Priority: 01.08.1989 US 388042

(74) Representative: Goodman, Christopher et al

Eric Potter & Clarkson

St. Mary's Court

St. Mary's Gate

Nottingham NG1 1LE (GB)

(43) Date of publication of application:

06.02.1991 Bulletin 1991/06

(56) References cited:

EP-A- 0 306 211

US-A- 4 757 442

(73) Proprietor: DIGITAL EQUIPMENT CORPORATION

Maynard, MA 01754 (US)

(72) Inventors:

• Bissett, Thomas D.

Derry, New Hampshire 03038 (US)

EP 0 411 805 B1

Note: Within nine months from the publication of the mention of the grant of the European patent, any person may give notice to the European Patent Office of opposition to the European patent granted. Notice of opposition shall be filed in a written reasoned statement. It shall not be deemed to have been filed until the opposition fee has been paid. (Art. 99(1) European Patent Convention).

## D escription

### I. BACKGROUND OF THE INVENTION

The present invention relates to the field of synchronizing multiple computer systems and, more particularly, to the synchronization of a dual-zone computer system in which the different computer system zones are designed to operate in lockstep with each other.

In a fault-tolerant computer system having two or more identical computer system zones, there needs to be a mechanism to synchronize the zones, either during start-up or during operation, to achieve a full duplex operation with both zones in lockstep. During operation, the need for synchronization can arise when one computer system zone has been disabled due to an error or when one zone undergoes routine maintenance.

One way to synchronize the zones is to stop both zones, and thus the entire computer system, to put the zones into the same state, and to then restart them. Although this technique may not be objectionable if the entire fault tolerant computer system is already down, such as during initialization or during a complete system failure, any resynchronization technique which requires disabling the entire computer system for a substantial period of time when only one zone is not operating has serious disadvantages.

Fault-tolerant computer systems are most often used in critical applications, such as air traffic control or control of nuclear processing plants, in which any down time could be extremely harmful or costly. The longer the down time, the greater the harm or potential harm.

Indeed, the whole rationale for using multiple zones is to avoid disabling an entire system when only one zone is operating improperly. Thus, rather than bringing the entire computer system down for resynchronization, efforts should be made to keep the one good zone operating for as long a period of time as possible.

Because the fault tolerant system must be returned to full multiple zone capability if it is to remain fault tolerant, the inoperative computer system zone, once repaired, must be resynchronized to the operating zone. Such resynchronization is difficult for many reasons. One is the difficulty in determining when the two zones are in identical states and thus capable of being synchronized. Because the operating computer system zone is continually changing state, it is difficult to ensure that the non-operating zone is maintained in the same state as the operating zone.

Another reason for difficulty stems from the independence of the zones. In a computer system having two or more zones which operate independently, the zones each generate their own certain internal signals, such as memory refresh signals, to maintain proper operation. This allows each of the zones to operate independently should the other zone stop operating properly. During synchronization, however, this independent operation can prove to be troublesome because the in-

operative zone's internal signals, such as its memory refresh signals, may prevent that zone from properly synchronizing with the operating zone.

There is a need for a multiple zone computer system which allows on-line synchronization. Such synchronization should also attempt to minimize the disruption of the operating zone during synchronization.

US-4757442 describes a multi-processor system in which three processing systems, each including a CPU and corresponding main memory operate in synchronism with one another. When an asynchronism is detected, a faulting processor is identified by majority consensus and is isolated. To resynchronize, the processors must interrupt execution. Each processor is provided with a control memory in which a microprogram is stored. The contents of a master processor are copied to the asynchronous processor and all processors restart execution from a fixed address of the microprogram.

20

### II. SUMMARY OF THE INVENTION

The present invention overcomes the problems of synchronization by transferring to the memory system of a zone which is not operating, the contents of the memory system in the operating zone as well as all the storage commands received by the memory system of the operating zone during synchronization. Once full transfer of memory contents takes place, synchronization can occur by briefly resetting both zones into the same state and resuming the operating systems in both zones.

More specifically, in accordance with the invention, as embodied and as broadly described herein, the method of synchronizing a slave zone of a dual zone computer system to a master zone of the dual zone computer system is provided in accordance with claim 1.

A dual zone computer system, in accordance with the present invention as defined in claim 7, is capable of operating both in a synchronized mode and a synchronizing mode and comprises a master zone and a slave zone. The slave zone includes a slave zone CPU containing a plurality of CPU registers which define the state of the CPU, a slave zone memory system coupled to the slave zone CPU and being responsive to access commands and storage commands including write and memory refresh commands, from the slave zone CPU and from sources other than the slave zone CPU, and a slave zone cross-link means, coupled to the slave zone memory system, for receiving the storage commands and data from the master zone when the computer system is in the synchronizing mode. The master zone includes a master zone CPU also containing a plurality of CPU registers which define the state of the master zone CPU, a master zone memory system coupled to the master zone CPU and being responsive to access commands and storage commands, including write and memory refresh commands, both from the master zone

CPU and from sources other than the master zone CPU, and a master zone cross-link means for routing data stored in the master zone memory system and storage commands received by the master zone memory system to the slave zone when the computer system is in the synchronizing mode.

### III. BRIEF DESCRIPTION OF THE DRAWINGS

The accompanying drawings, which are incorporated in and which constitute a part of this specification, illustrate one embodiment of the invention and, together with the description of the invention, explain the principles of the invention.

Fig. 1 is a block diagram of a preferred embodiment of fault tolerant computer system which practices the present invention;

Fig. 2 is an illustration of the physical hardware containing the fault tolerant computer system in Fig. 1; Fig. 3. is a block diagram of the CPU module shown in the fault tolerant computer system shown in Fig. 1;

Fig. 4 is a block diagram of an interconnected CPU module and I/O module for the computer system shown in Fig. 1;

Fig. 5 is a block diagram of a memory module for the fault tolerant computer system shown in Fig. 1; Fig. 6 is a detailed diagram of the elements of the control logic in the memory module shown in Fig. 5; Fig. 7 is a block diagram of portions of the primary memory controller of the CPU module shown in Fig. 3;

Fig. 8 is a block diagram of the DMA engine in the primary memory controller of the CPU module of Fig. 3;

Fig. 9 is a diagram of error processing circuitry in the primary memory controller of the CPU module of Fig. 3;

Fig. 10 is a drawing of some of the registers of the cross-link in the CPU module shown in Fig. 3;

Fig. 11 is a block diagram of the elements which route control signals in the cross-links of the CPU module shown in Fig. 3;

Fig. 12 is a block diagram of the elements which route data and address signals in the primary cross-link of the CPU module shown in Fig. 3;

Fig. 13 is a state diagram showing the states for the cross-link of the CPU module shown in Fig. 3;

Fig. 14 is a block diagram of the timing system for the fault tolerant computer system of Fig. 1;

Fig. 15 is a timing diagram for the clock signals generated by the timing system in Fig. 14;

Fig. 16 is a detailed diagram of a phase detector for the timing system shown in Fig. 14;

Fig. 17 is a block diagram of an I/O module for the computer system of Fig. 1;

Fig. 18 is a block diagram of the firewall element in

the I/O module shown in Fig. 17;

Fig. 19 is a detailed diagram of the elements of the cross-link pathway for the computer system of Fig. 1;

Figs. 20A-20E are data flow diagrams for the computer system in Fig. 1;

Fig. 21 is a block diagram of zone 20 showing the routing of reset signals;

Fig. 22 is a block diagram of the components involved in resets in the CPU module shown in Fig. 3;

Fig. 23 is a diagram of clock reset circuitry; and Fig. 24 is logic flow diagram for resynchronization and bulk memory transfer by the computer system in Fig. 1.

### IV. DESCRIPTION OF THE PREFERRED EMBODIMENT

Reference will now be made in detail to a presently preferred embodiment of the invention, an example of which is illustrated in the accompanying drawings.

#### A. SYSTEM DESCRIPTION

Fig. 1 is a block diagram of a fault tolerant computer system 10 in accordance with the present invention.

Fault tolerant computer system 10 includes duplicate systems, called zones. In the normal mode, the two zones 11 and 11' operate simultaneously. The duplication ensures that there is no single point of failure and that a single error or fault in one of the zones 11 or 11' will not disable computer system 10. Furthermore, all such faults can be corrected by disabling or ignoring the device or element which caused the fault. Zones 11 and 11' are shown in Fig. 1 as respectively including duplicate processing systems 20 and 20'. The duality, however, goes beyond the processing system.

Fig. 2 contains an illustration of the physical hardware of fault tolerant computer system 10 and graphically illustrates the duplication of the systems. Each zone 11 and 11' is housed in a different cabinet 12 and 12', respectively. Cabinet 12 includes battery 13, poor regulator 14, cooling fans 16, and AC input 17. Cabinet 12' includes separate elements corresponding to elements 13, 14, 16 and 17 of cabinet 12.

As explained in greater detail below, processing systems 20 and 20' include several modules interconnected by backplanes. If a module contains a fault or error, that module may be removed and replaced without disabling computing system 10. This is because processing systems 20 and 20' are physically separate, have separate backplanes into which the modules are plugged, and can operate independently of each other. Thus modules can be removed from and plugged into the backplane of one processing system while the other processing system continues to operate.

In the preferred embodiment, the duplicate processing systems 20 and 20' are identical and contain

identical modules. Thus, only processing system 20 will be described completely with the understanding that processing system 20' operates equivalently.

Processing system 20 includes CPU module 30 which is shown in greater detail in Figs. 3 and 4. CPU module 30 is interconnected with CPU module 30' in processing system 20' by a cross-link pathway 25 which is described in greater detail below. Cross-link pathway 25 provides data transmission paths between processing systems 20 and 20' and carries timing signals to ensure that processing systems 20 and 20' operate synchronously.

Processing system 20 also includes I/O modules 100, 110, and 120. I/O modules 100, 110, 120, 100', 110' and 120' are independent devices. I/O module 100 is shown in greater detail in Figs. 1, 4, and 17. Although multiple I/O modules are shown, duplication of such modules is not a requirement of the system. Without such duplication, however, some degree of fault tolerance will be lost.

Each of the I/O modules 100, 110 and 120 is connected to CPU module 30 by dual rail module interconnects 130 and 132. Module interconnects 130 and 132 serve as the I/O interconnect and are routed across the backplane for processing system 20. For purposes of this application, the data pathway including CPU 40, memory controller 70, cross-link 90 and module interconnect 130 is considered as one rail, and the data pathway including CPU 50, memory controller 75, cross-link 95, and module interconnect 132 is considered as another rail. During proper operation, the data on both rails is the same.

## B. FAULT TOLERANT SYSTEM PHILOSOPHY

Fault tolerant computer system 10 does not have a single point of failure because each element is duplicated. Processing systems 20 and 20' are each a fail stop processing system which means that those systems can detect faults or errors in the subsystems and prevent uncontrolled propagation of such faults and errors to other subsystems, but they have a single point of failure because the elements in each processing system are not duplicated.

The two fail stop processing systems 20 and 20' are interconnected by certain elements operating in a defined manner to form a fail safe system. In the fail safe system embodied as fault tolerant computer system 10, the entire computer system can continue processing even if one of the fail stop processing systems 20 and 20' is faulting.

The two fail stop processing systems 20 and 20' are considered to operate in lockstep synchronism because CPUs 40, 50, 40' and 50' operate in such synchronism. There are three significant exceptions. The first is at initialization when a bootstrapping technique brings both processors into synchronism. The second exception is when the processing systems 20 and 20' operate inde-

pendently (asynchronously) on two different workloads. The third exception occurs when certain errors arise in processing systems 20 and 20'. In this last exception, the CPU and memory elements in one of the processing systems is disabled, thereby ending synchronous operation.

When the system is running in lockstep I/O, only one I/O device is being accessed at any one time. All four CPUs 40, 50, 40' and 50', however, would receive the same data from that I/O device at substantially the same time. In the following discussion, it will be understood that lockstep synchronization of processing systems means that only one I/O module is being accessed.

The synchronism of duplicate processing systems

20 and 20' is implemented by treating each system as a deterministic machine which, starting in the same known state and upon receipt of the same inputs, will always enter the same machine states and produce the same results in the absence of error. Processing systems 20 and 20' are configured identically, receive the same inputs, and therefore pass through the same states. Thus, as long as both processors operate synchronously, they should produce the same results and enter the same state. If the processing systems are not in the same state or produce different results, it is assumed that one of the processing systems 20 and 20' has faulted. The source of the fault must then be isolated in order to take corrective action, such as disabling the faulting module.

30 Error detection generally involves overhead in the form of additional processing time or logic. To minimize such overhead, a system should check for errors as infrequently as possible consistent with fault tolerant operation. At the very least, error checking must occur before data is outputted from CPU modules 30 and 30'.

35 Otherwise, internal processing errors may cause improper operation in external systems, like a nuclear reactor, which is the condition that fault tolerant systems are designed to prevent.

40 There are reasons for additional error checking. For example, to isolate faults or errors it is desirable to check the data received by CPU modules 30 and 30' prior to storage or use. Otherwise, when erroneous stored data is later accessed and additional errors result, it becomes difficult or impossible to find the original source of errors, especially when the erroneous data has been stored for some time. The passage of time as well as subsequent processing of the erroneous data may destroy any trail back to the source of the error.

50 "Error latency," which refers to the amount of time an error is stored prior to detection, may cause later problems as well. For example, a seldom-used routine may uncover a latent error when the computer system is already operating with diminished capacity due to a previous error. When the computer system has diminished capacity, the latent error may cause the system to crash.

55 Furthermore, it is desirable in the dual rail systems

of processing systems 20 and 20' to check for errors prior to transferring data to single rail systems, such as a shared resource like memory. This is because there are no longer two independent sources of data after such transfers, and if any error in the single rail system is later detected, then error tracing becomes difficult if not impossible. The preferred method of error handling is set forth in an application filed this same date entitled, "Software Error Handling", number 90308000.0, published is EP-0 415 545 A.

### C. MODULE DESCRIPTION

#### 1. CPU Module

The elements of CPU module 30 which appear in Fig. 1 are shown in greater detail in Figs. 3 and 4. Fig. 3 is a block diagram of the CPU module, and Fig. 4 shows block diagrams of CPU module 30 and I/O module 100 as well as their interconnections. Only CPU module 30 will be described since the operation of and the elements included in CPU modules 30 and 30' are generally the same.

CPU module 30 contains dual CPUs 40 and 50. CPUs 40 and 50 can be standard central processing units known to persons of ordinary skill. In the preferred embodiment, CPUs 40 and 50 are VAX microprocessors manufactured by Digital Equipment Corporation, the assignee of this application.

Associated with CPUs 40 and 50 are cache memories 42 and 52, respectively, which are standard cache RAMs of sufficient memory size for the CPUs. In the preferred embodiment, the cache RAM is 4K x 64 bits. It is not necessary for the present invention to have a cache RAM, however.

#### 2. Memory Module

Preferably, CPU's 40 and 50 can share up to four memory modules 60. Fig. 5 is a block diagram of one memory module 60 shown connected to CPU module 30.

During memory transfer cycles, status register transfer cycles, and EEPROM transfer cycles, each memory module 60 transfers data to and from primary memory controller 70 via a bidirectional data bus 85. Each memory module 60 also receives address, control, timing, and ECC signals from memory controllers 70 and 75 via buses 80 and 82, respectively. The address signals on buses 80 and 82 include board, bank, and row and column address signals that identify the memory board, bank, and row and column address involved in the data transfer.

As shown in Fig. 5, each memory module 60 includes a memory array 600. Each memory array 600 is a standard RAM in which the DRAMs are organized into eight banks of memory. In the preferred embodiment, fast page mode type DRAMs are used.

Memory module 60 also includes control logic 610, data transceivers/registers 620, memory drivers 630, and an EEPROM 640. Data transceivers/receivers 620 provide a data buffer and data interface for transferring data between memory array 600 and the bidirectional data lines of data bus 85. Memory drivers 630 distribute row and column address signals and control signals from control logic 610 to each bank in memory array 600 to enable transfer of a longword of data and its corresponding ECC signals to or from the memory bank selected by the memory board and bank address signals.

EEPROM 640, which can be any type of NVRAM (nonvolatile RAM), stores memory error data for off-line repair and configuration data, such as module size.

When the memory module is removed after a fault, stored data is extracted from EEPROM 640 to determine the cause of the fault. EEPROM 640 is addressed via row address lines from drivers 630 and by EEPROM control signals from control logic 610. EEPROM 640 transfers eight bits of data to and from a thirty-two bit internal memory data bus 645.

Control logic 610 routes address signals to the elements of memory module 60 and generates internal timing and control signals. As shown in greater detail in Fig. 6, control logic 610 includes a primary/mirror designator circuit 612.

Primary/mirror designator circuit 612 receives two sets of memory board address, bank address, row and column address, cycle type, and cycle timing signals from memory controllers 70 and 75 on buses 80 and 82, and also transfers two sets of ECC signals to or from the memory controllers on buses 80 and 82. Transceivers/registers in designator 612 provide a buffer and interface for transferring these signals to and from memory buses 80 and 82. A primary/mirror multiplexer bit stored in status registers 618 indicates which one of memory controllers 70 and 75 is designated as the primary memory controller and which is designated as the mirror memory controller, and a primary/mirror multiplexer signal is provided from status registers 618 to designator 612.

Primary/mirror designator 612 provides two sets of signals for distribution in control logic 610. One set of signals includes designated primary memory board address, bank address, row and column address, cycle type, cycle timing, and ECC signals. The other set of signals includes designated mirror memory board address, bank address, row and column address, cycle type, cycle timing, and ECC signals. The primary/mirror multiplexer signal is used by designator 612 to select whether the signals on buses 80 and 82 will be respectively routed to the lines for carrying designated primary signals and to the lines for carrying designated mirror signals, or vice-versa.

A number of time division multiplexed bidirectional lines are included in buses 80 and 82. At certain times after the beginning of memory transfer cycles, status register transfer cycles, and EEPROM transfer cycles,

ECC signals corresponding to data on data bus 85 are placed on these time division multiplexed bidirectional lines. If the transfer cycle is a write cycle, memory module 60 receives data and ECC signals from the memory controllers. If the transfer cycle is a read cycle, memory module 60 transmits data and ECC signals to the memory controllers. At other times during transfer cycles, address, control, and timing signals are received by memory module 60 on the time division multiplexed bidirectional lines. Preferably, at the beginning of memory transfer cycles, status register transfer cycles, and EEPROM transfer cycles, memory controllers 70 and 75 transmit memory board address, bank address, and cycle type signals on these timeshared lines to each memory module 60.

Preferably, row address signals and column address signals are multiplexed on the same row and column address lines during transfer cycles. First, a row address is provided to memory module 60 by the memory controllers, followed by a column address about sixty nanoseconds later.

A sequencer 616 receives as inputs a system clock signal and a reset signal from CPU module 30, and receives the designated primary cycle timing, designated primary cycle type, designated mirror cycle timing, and designated mirror cycle type signals from the transceivers/registers in designator 612.

Sequencer 616 is a ring counter with associated steering logic that generates and distributes a number of control and sequence timing signals for the memory module that are needed in order to execute the various types of cycles. The control and sequence timing signals are generated from the system clock signals, the designated primary cycle timing signals, and the designated primary cycle type signals.

Sequencer 616 also generates a duplicate set of sequence timing signals from the system clock signals, the designated mirror cycle timing signals, and the designated mirror cycle type signals. These duplicate sequence timing signals are used for error checking. For data transfers of multi-long words of data to and from memory module 60 in a fast page mode, each set of column addresses starting with the first set is followed by the next column address 120 nanoseconds later, and each long word of data is moved across bus 85 120 nanoseconds after the previous long word of data.

Sequencer 616 also generates tx/rx register control signals. The tx/rx register control signals are provided to control the operation of data transceivers/registers 620 and the transceivers/registers in designator 612. The direction of data flow is determined by the steering logic in sequencer 616, which responds to the designated primary cycle type signals by generating tx/rx control and sequence timing signals to indicate whether and when data and ECC signals should be written into or read from the transceivers/registers in memory module 60. Thus, during memory write cycles, status register write cycles, and EEPROM write cycles, data and ECC

signals will be latched into the transceivers/registers from buses 80, 82, and 85, while during memory read cycles, status register read cycles, and EEPROM read cycles, data and ECC signals will be latched into the transceivers/registers from memory array 600, status registers 618, or EEPROM 640 for output to CPU module 30.

Sequencer 616 also generates EEPROM control signals to control the operation of EEPROM 640.

The timing relationships that exist in memory module 60 are specified with reference to the rise time of the system clock signal, which has a period of thirty nanoseconds. All status register read and write cycles, and all memory read and write cycles of a single longword, are performed in ten system clock periods, i.e., 300 nanoseconds. Memory read and write transfer cycles may consist of multi-longword transfers. For each additional longword that is transferred, the memory transfer cycle is extended for four additional system clock periods. Memory refresh cycles and EEPROM write cycles require at least twelve system clock periods to execute, and EEPROM read cycles require at least twenty system clock periods.

The designated primary cycle timing signal causes sequencer 616 to start generating the sequence timing and control signals that enable the memory module selected by the memory board address signals to implement a requested cycle. The transition of the designated primary cycle timing signal to an active state marks the start of the cycle. The return of the designated primary cycle timing signal to an inactive state marks the end of the cycle.

The sequence timing signals generated by sequencer 616 are associated with the different states entered by the sequencer as a cycle requested by CPU module 30 is executed. In order to specify the timing relationship among these different states (and the timing relationship among sequence timing signals corresponding to each of these states), the discrete states that may be entered by sequencer 616 are identified as states SEQ IDLE and SEQ 1 to SEQ 19. Each state lasts for a single system clock period (thirty nanoseconds). Entry by sequencer 616 into each different state is triggered by the leading edge of the system clock signal. The leading edges of the system clock signal that cause sequencer 616 to enter states SEQ IDLE and SEQ 1 to SEQ 19 are referred to as transitions T IDLE and T1 to T19 to relate them to the sequencer states, i.e., TN is the system clock signal leading edge that causes sequencer 616 to enter state SEQ N.

At times when CPU module 30 is not directing memory module 60 to execute a cycle, the designated primary cycle timing signal is not asserted, and the sequencer remains in state SEQ IDLE. The sequencer is started (enters state SEQ 1) in response to assertion by memory controller 70 of the cycle timing signal on bus 80, provided control logic 610 and sequencer 616 are located in the memory module selected by memory

board address signals also transmitted from memory controller 70 on bus 80. The rising edge of the first system clock signal following assertion of the designated primary cycle active signal corresponds to transition T1.

As indicated previously, in the case of transfers of a single longword to or from memory array 600, the cycle is performed in ten system clock periods. The sequencer proceeds from SEQ IDLE, to states SEQ 1 through SEQ 9, and returns to SEQ IDLE.

Memory read and write cycles may be extended, however, to transfer additional longwords. Memory array 600 preferably uses "fast page mode" DRAMs. During multi-longword reads and writes, transfers of data to and from the memory array after transfer of the first longword are accomplished by repeatedly updating the column address and regenerating a CAS (column address strobe) signal.

During multi-longword transfer cycles, these updates of the column address can be implemented because sequencer 616 repeatedly loops from states SEQ 4 through SEQ 7 until all of the longwords are transferred. For example, if three longwords are being read from or written into memory array 600, the sequencer enters states SEQ IDLE, SEQ 1, SEQ 2, SEQ 3, SEQ 4, SEQ 5, SEQ 6, SEQ 7, SEQ 4, SEQ 5, SEQ 6, SEQ 7, SEQ 4, SEQ 5, SEQ 6, SEQ 7, SEQ 8, SEQ 9, and SEQ IDLE.

During a memory transfer cycle, the designated primary cycle timing signal is monitored by sequencer 616 during transition T6 to determine whether to extend the memory read or write cycle in order to transfer at least one additional longword. At times when the designated primary cycle timing signal is asserted during transition T6, the sequencer in state SEQ 7 will respond to the next system clock signal by entering state SEQ 4 instead of entering state SEQ 8.

In the case of a multi-longword transfer, the designated primary cycle timing signal is asserted at least fifteen nanoseconds before the first T6 transition and remains asserted until the final longword is transferred. In order to end a memory transfer cycle after the final longword has been transferred, the designated primary cycle timing signal is deasserted at least fifteen nanoseconds before the last T6 transition and remains deasserted for at least ten nanoseconds after the last T6 transition.

During memory transfer cycles, the designated primary row address signals and the designated primary column address signals are presented at different times by designator 612 in control logic 610 to memory drivers 630 on a set of time division multiplexed lines. The outputs of drivers 630 are applied to the address inputs of the DRAMs in memory array 600, and also are returned to control logic 610 for comparison with the designated mirror row and column address signals to check for errors. During status register transfer cycles and EEPROM transfer cycles, column address signals are not needed to select a particular storage location.

During a memory transfer cycle, row address signals are the first signals presented on the timeshared row and column address lines of buses 80 and 82. During state SEQ IDLE, row address signals are transmitted

5 by the memory controllers on the row and column address lines, and the row address is stable from at least fifteen nanoseconds before the T1 transition until ten nanoseconds after the T1 transition. Next, column address signals are transmitted by the memory controllers on the row and column address lines, and the column address is stable from at least ten nanoseconds before the T3 transition until fifteen nanoseconds after the T4 transition. In the case of multi-longword transfers during memory transfer cycles, subsequent column address signals

10 are then transmitted on the row and column address lines, and these subsequent column addresses are stable from ten nanoseconds before the T6 transition until fifteen nanoseconds after the T7 transition.

Generator/checker 617 receives the two sets of se-

20 quence timing signals generated by sequencer 616. In addition, the designated primary cycle type and bank address signals and the designated mirror cycle type and bank address signals are transmitted to generator/checker 617 by designator 612. In the generator/checker, a number of primary control signals, i.e., RAS (row address strobe), CAS (column address strobe) and WE (write enable), are generated for distribution to drivers 630, using the primary sequence timing signals and the designated primary cycle type and bank address signals.

25 A duplicate set of these control signals is generated by generator/checker 617 from the duplicate (mirror) sequence timing signals and the designated mirror cycle type and bank address signals. These mirror RAS, CAS, and write enable signals are used for error checking.

30 When the primary cycle type signals indicate a memory transfer cycle is being performed, the primary bank address signals identify one selected bank of DRAMs in memory array 600. Memory drivers 630 include separate RAS drivers for each bank of DRAMs in memory array 600. In generator/checker 617, the primary RAS signal is generated during the memory transfer cycle and demultiplexed onto one of the lines connecting the generator/checker to the RAS drivers. As a result, only the RAS driver corresponding to the selected

35 DRAM bank receives an asserted RAS signal during the memory transfer cycle. During refresh cycles, the primary RAS signal is not demultiplexed and an asserted RAS signal is received by each RAS driver. During status register transfer cycles and EEPROM transfer cycles, the bank address signals are unnecessary.

40 Memory drivers 630 also include CAS drivers. In generator/checker 617, the primary CAS signal is generated during memory transfer cycles and refresh cycles. The primary CAS signal is not demultiplexed and an asserted CAS signal is received by each CAS driver.

45 During memory write cycles, the primary WE signal is generated by generator/checker 617. The asserted WE signal is provided by drivers 630 to each DRAM

bank in memory array 600. However, a write can only be executed by the selected DRAM bank, which also receives asserted RAS and CAS signals.

In the preferred embodiment of the invention, during memory transfer cycles the primary signal is asserted during the T2 transition, is stable from at least ten nanoseconds before the T3 transition, and is deasserted during the last T7 transition. The primary CAS signal is asserted fifteen nanoseconds after each T4 transition, and is deasserted during each T7 transition. During memory write cycles the primary WE signal is asserted during the T3 transition, is stable from at least ten nanoseconds before the first T4 transition, and is deasserted during the last T7 transition.

When the primary cycle type signals indicate a memory refresh cycle is being performed, generator/checker 617 causes memory array 600 to perform memory refresh operations in response to the primary sequence timing signals provided by sequencer 616. During these refresh operations, the RAS and CAS signals are generated and distributed by the generator/checker in reverse order. This mode of refresh requires no external addressing for bank, row, or column.

During transfer cycles, ECC signals are transferred on the time division multiplexed bidirectional lines of buses 80 and 82 at times when data is being transferred on bus 85. However, these same lines are used to transfer control (e.g., cycle type) and address (e.g., memory board address and bank address) signals at other times during the transfer cycle.

The transceivers/registers in primary/mirror designator 612 include receivers and transmitters that are responsive to sequence timing signals and tx/rx register control signals provided by sequencer 616. The sequence timing signals and tx/rx register control signals enable multiplexing of ECC signals and address and control signals on the time division multiplexed bidirectional lines of buses 80 and 82.

Preferably, control and address signals, such as cycle type, memory board address, and bank address signals, are transmitted by memory controllers 70 and 75 and presented on the timeshared lines of buses 80 and 82 at the beginning of either single or multi-longword transfer cycles. These signals start their transition (while the sequencer is in the SEQ IDLE state) concurrent with activation of the cycle timing signal, and remain stable through T2. Therefore, in the transceivers/registers of designator 612, the receivers are enabled and the transmitters are set into their tristate mode at least until the end of state SEQ 2.

The cycle type signals identify which of the following listed functions will be performed by memory array 60 during the cycle: memory read, memory write, status register read, status register write, EEPROM read, EEPROM write, and refresh. The designated primary cycle type signals received by designator 612 are provided to sequencer 616 and used in generating tx/rx control signals and sequence timing signals. For example, in data

transceivers/registers 620 and in the transceivers/registers of designator 612, the receivers are enabled and the transmitters are set into their tristate mode by sequencer 616 throughout a write cycle. However, in data transceivers/registers 620 and in the transceivers/registers of designator 612 during a read cycle, the receivers are set into their tristate mode and the transmitters are enabled by sequencer 616 after the cycle type, memory board address, and bank address signals have been received at the beginning of the cycle.

In the preferred embodiment, data transferred to or from memory array 600 is checked in each memory module 60 using an Error Detecting Code (EDC), which is preferably the same code required by memory controllers 70 and 75. The preferred code is a single bit correcting, double bit detecting, error correcting code (ECC).

During a memory write cycle, memory controller 70 transmits at least one longword of data on data bus 85 and simultaneously transmits a corresponding set of ECC signals on bus 80. Meanwhile, memory controller 75 transmits a second set of ECC signals, which also correspond to the longword on data bus 85, on bus 82.

As embodied herein, during a memory write cycle the data and the ECC signals for each longword are presented to the receivers of data transceivers/registers 620 and to the receivers of the transceivers/registers of designator 612. The data and the ECC signals, which are stable at least ten nanoseconds before the T4 transition and remain stable until fifteen nanoseconds after the T6 transition, are latched into these transceivers/registers. During this time period, memory controllers 70 and 75 do not provide address and control signals on the timeshared lines of buses 80 and 82.

The designated primary ECC signals received by designator 612 and the longword of data received by transceivers/registers 620 during the memory write cycle are provided to the data inputs of the DRAMs in each of the eight banks of memory array 600 and to ECC generator 623. The generated ECC is cared to the designated primary ECC by comparator 625. The designated primary ECC signals also are provided to ECC comparators 625, together with the designated mirror ECC signals.

As embodied herein, during a memory read cycle, at least one longword of data and a corresponding set of ECC signals are read from memory array 600 and respectively steered to data transceivers/registers 620 and to the transceivers/registers of designator 612. During transition T7 of the memory read cycle, the data and the ECC signals for each longword are available from memory array 600 and are latched into these transceivers/registers. The data is also presented to the ECC generator 623 and its output is compared to the ECC read from memory.

After latching, the data and the ECC signals are presented to data bus 85 and to buses 80 and 82 by the transmitters of data transceivers/registers 620 and by

the transmitters of the transceivers/registers of designator 612. The same ECC signals are transmitted from the transceivers/registers in designator 612 to memory controller 70 and to memory controller 75. The data and the ECC signals transmitted on data bus 85 and on buses 80 and 82 are stable from fifteen nanoseconds after the T7 transition until five nanoseconds before the following T6 transition (in the case of a multi-longword transfer) or until five nanoseconds before the following T IDLE transition (in the case of a single longword transfer or the last longword of a multi-longword transfer). During this time period, memory controllers 70 and 75 do not provide address and control signals on the timeshared lines of buses 80 and 82. The transmitters of data transceivers/registers 620 and the transmitters of the transceivers/registers of designator 612 are set into their tristate mode during the following T IDLE transition.

Comparator 614 is provided to compare the address, control, and timing signals originating from controller 70 with the corresponding address, control, and timing signals originating from controller 75. The designated primary cycle timing signals, cycle type signals, memory board address signals, and bank address signals, together with the designated mirror cycle timing signals, cycle type signals, memory board address signals, bank address signals, row address signals, and column address signals, are provided from designator 612 to comparator 614. The designated primary row address signals and column address signals are provided from the outputs of drivers 630 to comparator 614. Both sets of signals are then compared.

If there is a miscompare between any of the address, control, and timing signals originating from the memory controllers, comparator 614 generates an appropriate error signal. As shown in Figure 6, board address error, bank address error, row address error, column address error, cycle type address error and cycle timing error signals may be output by the comparator.

Generator/checker 617 compares the primary control and timing signals generated by sequencer 616 and generator/checker 617 using the designated primary bank address, cycle type, and cycle timing signals with the mirror control and timing signals generated using the designated mirror bank address, cycle type, and cycle timing signals. The two sets of sequence timing signals are provided by sequencer 616 to generator/checker 617. The primary RAS, CAS, and WE signals are provided from the outputs of drivers 630 to generator/checker 617. As indicated previously, the mirror RAS, CAS, and WE signals are generated internally by the generator/checker. Generator/checker 617 compares the primary RAS, CAS, WE, and sequence timing signals to the mirror RAS, CAS, WE, and sequence timing signals.

If there is a miscompare between any of the control and timing signals originating from sequencer 616 or generator/checker 617, the generator/checker generates an appropriate error signal. As shown in Figure 6,

sequencer error, RAS error, CAS error, and WE error signals may be output by generator/checker 617.

Error signals are provided from comparator 614 and from generator/checker 617 to address/control error logic 621.

5 In response to receipt of an error signal from comparator 614 or from generator/checker 617, address/control error logic 621 transmits an address/control error signal to CPU module 30 to indicate the detection of a fault due to a miscompare between any address, control, or timing signals. The address/control error signal is sent to error logic in memory controllers 70 and 75 for error handling. The transmission of the address/control error signal to CPU module 30 causes a CPU/MEM fault, which is discussed in greater detail in other sections.

10 The error signals from comparator 614 and from generator/checker 617 also are provided to status registers 618. In the status registers, the error signals and all of the address, control, timing, data, and ECC signals 15 relevant to the fault are temporarily stored to enable error diagnosis and recovery.

15 In accordance with one aspect of the invention, only a single thirty-two bit data bus 85 is provided between CPU module 30 and memory module 60. Therefore, 20 memory module 60 cannot compare two sets of data from memory controllers 70 and 75. However, data integrity is verified by memory module 60 without using a duplicate set of thirty-two data lines by checking the two 25 separate sets of ECC signals that are transmitted by memory controllers 70 and 75 to memory module 60.

30 As shown in Fig. 6, control logic 610 includes ECC generator 623 and ECC comparators 625. The designated primary and mirror ECC signals are provided by designator 612 to the ECC comparators. During a memory write cycle, the designated primary ECC signals are compared to the designated mirror ECC signals. As a result, memory module 60 verifies whether memory controllers 70 and 75 are in agreement and whether the 35 designated primary ECC signals being stored in the DRAMs of memory array 600 during the memory write cycle are correct. Furthermore, the data presented to the data inputs of the DRAMs during the memory write cycle is provided to ECC generator 623. ECC generator 40 623 produces a set of generated ECC signals that correspond to the data and provides the generated ECC signals to ECC comparators 625. The designated primary ECC signals are compared to the generated ECC signals to verify whether the data transmitted on data bus 85 by memory controller 70 is the same as the data 45 being stored in the DRAMs of memory array 600.

50 During a memory read cycle, the data read from the selected bank of DRAMs is presented to the ECC generator. The generated ECC signals then are provided to the ECC comparators, which also receive stored ECC signals read from the selected bank of DRAMs. The generated and stored ECC signals are compared by ECC comparators 625.

If there is a miscompare between any of pairs of

ECC signals monitored by ECC comparators 625, the ECC comparators generate an appropriate error signal. As shown in Figure 6, primary/mirror ECC error, primary/generated ECC error, and memory/generated ECC error signals may be output by the ECC comparators.

These ECC error signals from ECC comparators 625 are provided to status registers 618. In the status registers, each of the ECC error signals and all of the address, control, timing, data, and ECC signals relevant to an ECC fault are temporarily stored to enable error diagnosis and recovery.

An ECC error signal is asserted by ECC comparators 625 on an ECC error line and transmitted to CPU module 30 to indicate the detection of an ECC fault due to a miscompare. The miscompare can occur during either of the two ECC check performed during a memory write cycle, or during the single check performed during a memory read cycle.

As shown in Figure 6, board select logic 627 receives slot signals from memory backplane. The slot signals specify a unique slot location for each memory module 60. Board select logic 627 then compares the slot signals with the designated primary board address signals transmitted from one of the memory controllers via designator circuit 612. A board selected signal is generated by board select logic 627 if the slot signals are the same as the designated primary board address signals, thereby enabling the other circuitry in control logic 610.

### 3. Memory Controller

Memory controllers 70 and 75 control the access of CPUs 40 and 50, respectively, to memory module 60, auxiliary memory elements and, in the preferred embodiment, perform certain error handling operations. The auxiliary memory elements coupled to memory controller 70 include system ROM 43, EEPROM 44, and scratch pad RAM 45. ROM 43 holds certain standard code, such as diagnostics, console drivers, and part of the bootstrap code. EEPROM 44 is used to hold information such as error information detected during the operation of CPU 40, which may need to be modified, but which should not be lost when power is removed. Scratch pad RAM 45 is used for certain operations performed by CPU 40 and to convert rail-unique information (e.g., information specific to conditions on one rail which is available to only one CPU 40 or 50) to zone information (e.g., information which can be accessed by both CPUs 40 and 50).

Equivalent elements 53, 54 and 55 are coupled to memory controller 75. System ROM 53, EEPROM 54, and scratch pad RAM 55 are the same as system ROM 43, EEPROM 44, and scratch pad RAM 45, respectively, and perform the same functions.

The details of the preferred embodiment of primary memory controller 70 can be seen in Figs. 7-9. Mirror memory controller 75 has the same elements as shown

in Figs. 7-9, but differs slightly in operation. Therefore, only primary memory controller 70's operation will be described, except where the operation of memory controller 75 differs. Memory controllers 70' and 75' in processing system 20' have the same elements and act the same as memory controllers 70 and 75, respectively.

These ECC error signals from ECC comparators 625 are provided to status registers 618. In the status registers, each of the ECC error signals and all of the address, control, timing, data, and ECC signals relevant to an ECC fault are temporarily stored to enable error diagnosis and recovery.

The elements shown in Fig. 7 control the flow of data, addresses and signals through primary memory controller 70. Control logic 700 controls the state of the various elements in Fig. 7 according to the signals received by memory controller 70 and the state engine of that memory controller which is stored in control logic 700. Multiplexer 702 selects addresses from one of three sources. The addresses can either come from CPU 30

via receiver 705, from the DMA engine 800 described below in reference to Fig. 8, or from a refresh resync address line which is used to generate an artificial refresh during certain bulk memory transfers from one zone to another during resynchronization operations.

The output of multiplexer 702 is an input to multiplexer 710, as is data from CPU 30 received via receiver 705 and data from DMA engine 800. The output of multiplexer 710 provides data to memory module 60 via memory interconnect 85 and driver 715. Driver 715 is disabled for mirror memory control modules 75 and 75' because only one set of memory data is sent to memory modules 60 and 60', respectively.

The data sent to memory interconnect 85 includes either data to be stored in memory module 60 from CPU 30 or DMA engine 800. Data from CPU 30 and addresses from multiplexer 702 are also sent to DMA engine 800 via this path and also via receiver 745 and ECC corrector 750.

The addresses from multiplexer 702 also provide an input to demultiplexer 720 which divides the addresses into a row/column address portion, a board/bank address portion, and a single board bit. The twenty-two bits of the row/column address are multiplexed onto eleven lines. In the preferred embodiment, the twenty-two row/column address bits are sent to memory module 60 via drivers 721. The single board bit is preferably sent to memory module 60 via driver 722, and the other board/bank address bits are multiplexed with ECC signals.

Multiplexer 725 combines a normal refresh command for memory controller 70 along with cycle type information from CPU 30 (i.e., read, write, etc.) and DMA cycle type information. The normal refresh command and the refresh resync address both cause memory module 60 to initiate a memory refresh operation.

The output of multiplexer 725 is an input to multiplexer 730 along with the board/bank address from demultiplexer 720. Another input into multiplexer 730 is the output of ECC generator/checker 735. Multiplexer 730 selects one of the inputs and places it on the time-division multiplexed ECC/address lines to memory module 60. Multiplexer 730 allows those time-division multiplexed lines to carry board/bank address and additional control information as well as ECC information, although

at different times.

ECC information is received from memory modules 60 via receiver 734 and is provided as an input to ECC generator/checker 735 to compare the ECC generated by memory module 60 with that generated by memory controller 70.

Another input into ECC generator/checker 735 is the output of multiplexer 740. Depending upon whether the memory transaction is a write transaction or a read transaction, multiplexer 740 receives as inputs the memory data sent to memory module 60 from multiplexer 710 or the memory data received from memory module 60 via receiver 745. Multiplexer 740 selects one of these sets of memory data to be the input to ECC generator/checker 735. Generator/checker 735 then generates the appropriate ECC code which, in addition to being sent to multiplexer 730, is also sent to ECC corrector 750. In the preferred embodiment, ECC corrector 750 corrects any single bit errors in the memory data received from memory module 60.

The corrected memory data from ECC checker 750 is then sent to the DMA engine shown in Fig. 8 as well as to multiplexer 752. The other input into multiplexer 752 is error information from the error handling logic described below in connection with Fig. 9. The output of multiplexer 752 is sent to CPU 30 via driver 753.

Comparator 755 compares the data sent from multiplexer 710 to memory module 60 with a copy of that data after it passes through driver 715 and receiver 745. This checking determines whether driver 715 and receiver 745 are operating correctly. The output of comparator 755 is a CMP error signal which indicates the presence or absence of such a comparison error. The CMP error feeds the error logic in Fig. 9.

Two other elements in Fig. 7 provide a different kind of error detection. Element 760 is a parity generator. ECC data, generated either by the memory controller 70 on data to be stored in memory module 60 or generated by memory module 60 on data read from memory module 60 is sent to a parity generator 760. The parity signal from generator 760 is sent, via driver 762, to comparator 765. Comparator 765 compares the ECC parity signal from generator 760 with an equivalent ECC parity signal generated by controller 75'.

Parity generator 770 performs the same type of a check on the row/column and single bit board address signals received from demultiplexer 720. The address parity signal from parity generator 770 is transmitted by a driver 772 to a comparator 775 which also receives an address parity signal from controller 75. The outputs of comparator 765 and 775 are parity error signals which feed the error logic in Fig. 9.

Fig. 8 shows the fundamentals of a DMA engine 800. In the preferred embodiment, DMA engine 800 resides in memory controller 70, but there is no requirement for such placement. As shown in Fig. 8, DMA engine 800 includes a data router 810, a DMA control 820, and DMA registers 830. Driver 815 and receiver 816

provide an interface between memory controller 70 and cross-link 90.

5 DMA control 820 receives internal control signals from control logic 700 and, in response, sends control signals to place data router 810 into the appropriate configuration. Control 820 also causes data router 810 to set its configuration to route data and control signals from cross-link 90 to the memory control 70 circuitry shown in Fig. 7. Data router 810 sends its status signals to DMA control 820 which relays such signals, along with other DMA information, to error logic in Fig. 9.

Registers 830 includes a DMA byte counter register 832 and a DMA address register 836. These registers are set to initial values by CPU 40 via router 810. Then, 15 during DMA cycles, control 820 causes, via router 810, the counter register 832 to increment and address register 836 to decrement. Control 820 also causes the contents of address registers 836 to be sent to memory module 60 through router 810 and the circuitry in Fig. 7 during DMA operations.

20 As explained above, in the preferred embodiment of this invention, the memory controllers 70, 75, 70' and 75' also perform certain fundamental error operations. An example of the preferred embodiment of the hardware to perform such error operations are shown in Fig. 9.

25 As shown in Fig. 9, certain memory controller internal signals, such as timeout, ECC error and bus miscompare, are inputs into diagnostic error logic 870, as 30 are certain external signals such as rail error, firewall miscompare, and address/control error. In the preferred embodiment, diagnostic error logic 870 receives error signals from the other components of system 10 via cross-links 90 and 95.

35 Diagnostic error logic 870 forms error pulses from the error signals and from a control pulse signal generated from the basic timing of memory controller 70. The error pulses generated by diagnostic error logic 870 contain certain error information which is stored into appropriate locations in a diagnostic error register 880 in accordance with certain timing signals. System fault error address register 865 stores the address in memory module 60 which CPUs 40 and 50 were communicating with when an error occurred.

40 45 The error pulses from diagnostic error logic 870 are also sent to error categorization logic 850 which also receives information from CPU 30 indicating the cycle type (e.g., read, write, etc.). From that information and the error pulses, error categorization logic 850 determines the presence of CPU/IO errors, DMA errors, or CPU/MEM faults.

50 55 A CPU/IO error is an error on an operation that is directly attributable to a CPU/IO cycle on bus 46 and may be hardware recoverable, as explained below in regard to resets. DMA errors are errors that occur during a DMA cycle and, in the preferred embodiment, are handled principally by software. CPU/MEM faults are errors that for which the correct operation of CPU or the con-

tents of memory cannot be guaranteed.

The outputs from error categorization logic 850 are sent to encoder 855 which forms a specific error code. This error code is then sent to cross-links 90 and 95 via AND gate 856 when the error disable signal is not present.

After receiving the error codes, cross-links 90, 95, 90' and 95' send a retry request signal back to the memory controllers. As shown in Fig. 9, an encoder 895 in memory controller 70 receives the retry request signal along with cycle type information and the error signals (collectively shown as cycle qualifiers). Encoder 895 then generates an appropriate error code for storage in a system fault error register 898.

System fault error register 898 does not store the same information as diagnostic error register 880. Unlike the system fault error register 898, the diagnostic error register 880 only contains rail unique information, such as an error on one input from a cross-link rail, and zone unique data, such as an uncorrectable ECC error in memory module 60.

System fault error register 898 also contains several bits which are used for error handling. These include a NXM bit indicating that a desired memory location is missing, a NXIO bit indicating that a desired I/O location is missing, a solid fault bit and a transient bit. The transient and solid bits together indicate the fault level. The transient bit also causes system fault error address register 865 to freeze.

Memory controller status register 875, although technically not part of the error logic, is shown in Fig. 9 also. Register 875 stores certain status information such as a DMA ratio code in DMA ratio portion 877, an error disable code in error disable portion 878, and a mirror bus driver enable code in mirror bus driver enable portion 876. The DMA ratio code specifies the fraction of memory bandwidth which can be allotted to DMA. The error disable code provides a signal for disabling AND gate 856 and thus the error code. The mirror bus driver enable code provides a signal for enabling the mirror bus drivers for certain data transactions.

#### 4. Cross-link

Data for memory resync, DMA and I/O operations pass through cross-links 90 and 95. Generally, cross-links 90 and 95 provide communications between CPU module 30, CPO module 30' I/O modules 100, 110, 120, and I/O modules 100', 110', 120' (see Fig. 1).

Cross-links 90 and 95 contain both parallel registers 910 and serial registers 920 as shown in Fig. 10. Both types of registers are used for interprocessor communication in the preferred embodiment of this invention. During normal operation, processing systems 20 and 20' are synchronized and data is exchanged in parallel between processing systems 20 and 20' using parallel registers 910 in cross-links 90/95 and 90'/95', respectively. When processing systems 20 and 20' are not syn-

chronized, most notably during bootstrapping, data is exchanged between cross-links by way of serial registers 920.

The addresses of the parallel registers are in I/O space as opposed to memory space. Memory space refers to locations in memory module 60. I/O space refers to locations such as I/O and internal system registers, which are not in memory module 60.

Within I/O space, addresses can either be in system address space or zone address space. The term "system address space" refers to addresses that are accessible throughout the entire system 10, and thus by both processing systems 20 and 20'. The term "zone address space" refers to addresses which are accessible only by the zone containing the particular cross-link.

The parallel registers shown in Fig. 10 include a communications register 906 and an I/O reset register 908. Communications register 906 contains unique data to be exchanged between zones. Such data is usually zone-unique, such as a memory soft error (it is almost beyond the realm of probability that memory modules 60 and 60' would independently experience the same error at the same time).

Because the data to be stored into register 906 is unique, the address of communications register 906 for purposes of writing must be in zone address space. Otherwise, processing systems 20 and 20', because they are in lockstep synchronization and executing the same series of instruction at substantially the same time, could not store zone unique data into only the communications registers 906 in zone 11; they would have to store that same data into the communications registers 906' (not shown) in zone 11'.

The address of communications register 906 for reading, however, is in system address space. Thus, during synchronous operation, both zones can simultaneously read the communications register from one zone and then simultaneously read the communications register from the other zone.

I/O reset register 908 resides in system address space. The I/O reset register includes one bit per I/O module to indicate whether the corresponding module is in a reset state. When an I/O module is in a reset state, it is effectively disabled.

Parallel registers 910 also include other registers, but an understanding of those other registers is not necessary to an understanding of the present invention.

All of the serial cross-link registers 920 are in the zone specific space since they are used either for asynchronous communication or contain only zone specific information. The purpose of the serial cross-link registers and the serial cross-link is to allow processors 20 and 20' to communicate even though they are not running in lockstep synchronization (i.e., phase-locked clocks and same memory states). In the preferred embodiment, there are several serial registers, but they need not be described to understand this invention.

Control and status register 912 is a serial register

which contains status and control flags. One of the flags is an OSR bit 913 which is used for bootstrapping and indicates whether the processing system in the corresponding zone has already begun its bootstrapping process or whether the operating system for that zone is currently running, either because its bootstrapping process has completed, or because it underwent a resynchronization.

Control and status register 912 also contain the mode bits 914 for identifying the current mode of cross-link 90 and thus of processing system 20. Preferably mode bits include resync mode bits 915 and cross-link mode bits 916. Resync mode bits 915 identify cross-link 90 as being either in resync slave or resync master mode. The cross-link mode bits 916 identify cross-link 90 as being either in cross-link off, duplex, cross-link master, or cross-link slave mode.

One of the uses for the serial registers is a status read operation which allows the cross-link in one zone to read the status of the other zone's cross-link. Setting a status read request flag 918 in serial control and status register 912 sends a request for status information to cross-link 90'. Upon receipt of this message, cross-link 90' sends the contents of its serial control and status register 912' back to cross-link 90.

Fig. 11 shows some of the elements for routing control and status signals (referred to as "control codes") in primary cross-link 90 and mirror cross-link 95. Corresponding cross-link elements exist in the preferred embodiment within cross-links 90' and 95'. These codes are sent between the memory controllers 70 and 75 and the I/O modules coupled to module interconnects 130, 132, 130' and 132'.

Fig. 12 shows the elements in the preferred embodiment of primary cross-link 90 which are used for routing data and address signals. Corresponding cross-link elements exist in cross-links 95, 90' and 95'.

In Fig. 11, the elements for both the primary cross-link 90 and mirror cross-link 95 in processing system 20 are shown, although the hardware is identical, because of an important interconnection between the elements. The circuit elements in mirror cross-link 95 which are equivalent to elements in primary cross-link 90 are shown by the same number, except in the mirror controller the letter "m" is placed after the number.

With reference to Figs. 11 and 12, the elements include latches, multiplexers, drivers and receivers. Some of the latches, such as latches 933 and 933m, act as delay elements to ensure the proper timing through the cross-links and thereby maintain synchronization. As shown in Fig. 11, control codes from memory controller 70 are sent via bus 88 to latch 931 and then to latch 932. The reason for such latching is to provide appropriate delays to ensure that data from memory controller 70 passes through cross-link 90 simultaneously with data from memory controller 70'.

If codes from memory controller 70 are to be sent to processing system 20' via cross-link 90', then driver

937 is enabled. The control codes from memory controller 70 also pass through latch 933 and into multiplexer CSMUXA 935. If control codes are received into primary cross-link 90 from cross-link 90', then their path is through receiver 936 into latch 938 and also into multiplexer 935.

Control codes to multiplexer 935 determine the source of data, that is either from memory controller 70 or from memory controller 70', and place those codes on the output of multiplexer 935. That output is stored in latch 939, again for proper delay purposes, and driver 940 is enabled if the codes are to be sent to module interconnect 130.

The path for data and address signals, as shown in Fig. 12 is somewhat similar to the path of control signals shown in Fig. 11. The differences reflect the fact that during any one transaction, data and addresses are flowing in only one direction through cross-links 90 and 95, but control signals can be flowing in both directions during that transaction. For that same reason the data lines in busses 88 and 89 are bidirectional, but the control codes are not.

Data and addresses from the memory controller 70, via bus 88, enter latch 961, then latch 962, and then latch 964. As in Fig. 11, the latches in Fig. 12 provide proper timing to maintain synchronization. Data from memory controller 70' is buffered by receiver 986, stored in latch 988, and then routed to the input of multiplexer 966. The output of multiplexer 966 is stored in latch 968 and, if driver 969 is enabled, is sent to module interconnect 130.

The path for control codes to be sent to memory controller 70 is shown in Fig. 11. Codes from module interconnect 130 are first stored in latch 941 and then presented to multiplexer CSMUXC 942. Multiplexer 942 also receives control codes from parallel cross-link registers 910 and selects either the parallel register codes or the codes from latch 941 for transmission to latch 943. If those control codes are to be transmitted to cross-link 90', then driver 946 is enabled. Control codes from cross-link 90' (and thus from memory controller 70') are buffered by receiver 947, stored in latch 948, and presented as an input to multiplexer CSMUXD 945. CSMUXD 945 also receives as an input the output of latch 944 which stores the contents of latch 943.

Multiplexer 945 selects either the codes from module interconnect 130 or from cross-link 90' and presents those signals as an input to multiplexer CSMUXE 949. Multiplexer 949 also receives as inputs a code from the decode logic 970 (for bulk memory transfers that occur during resynchronization), codes from the serial cross-link registers 920, or a predetermined error code ERR. Multiplexer 949 then selects ones of those inputs, under the appropriate control, for storage in latch 950. If those codes are to be sent to memory controller 70, then driver 951 is activated.

The purpose of the error code ERR, which is an input into multiplexer 949, is to ensure that an error in one

of the rails will not cause the CPUs in the same zone as the rails to process different information. If this occurred, CPU module 30 would detect a fault which would cause drastic, and perhaps unnecessary action. To avoid this, cross-link 90 contains an EXCLUSIVE OR gate 960 which compares the outputs of multiplexers 945 and 945m. If they differ, then gate 960 causes multiplexer 949 to select the ERR code. EXCLUSIVE OR gate 960m similarly causes multiplexer 949m also to select an ERR code. This code indicates to memory controllers 70 and 75 that there has been an error, but avoids causing a CPU module error. The single rail interface to memory module 60 accomplishes the same result for data and addresses.

The data and address flow shown in Fig. 12 is similar to the flow of control signals in Fig. 11. Data and addresses from module interconnect 130 are stored in latch 972 and then provided as an input to multiplexer MUXB 974. Data from the parallel registers 910 provide another input to multiplexer 974. The output of multiplexer 974 is an input to multiplexer MUXC 976 which also receives data and addresses stored in latch 961 that were originally sent from memory controller 70. Multiplexer 976 then selects one of the inputs for storage in latch 978. If the data and addresses, either from the module interconnect 130 or from the memory controller 70, are to be sent to cross-link 90', then driver 984 is enabled.

Data from cross-link 90' is buffered by receiver 986 and stored in latch 988, which also provides an input to multiplexer MUXD 982. The other input of multiplexer MUXD 982 is the output of latch 980 which contains data and addresses from latch 978. Multiplexer 982 then selects one of its inputs which is then stored into latch 990. If the data or addresses are to be sent to memory controller 70, then driver 992 is activated. Data from serial registers 920 are sent to memory controller 70 via driver 994.

The data routing in cross-link 90, and more particularly the cross-link elements in both Figs. 11 and 12, is controlled by several signals generated by decode logic 970, decode logic 971, decode logic 996, and decode logic 998. This logic provides the signals which control multiplexers 935, 942, 945, 949, 966, 974, 976, and 982 to select the appropriate input source. In addition, the decode logic also controls drivers 940, 946, 951, 969, 984, 992, and 994.

Most of the control signals are generated by decode logic 998, but some are generated by decode logic 970, 971, 970m, 971m, and 996. Decode logic 998, 970 and 970m are connected at positions that will ensure that the logic will receive the data and codes necessary for control whether the data and codes are received from its own zone or from other zone.

The purpose of decode logic 971, 971m and 996 is to ensure that the drivers 937, 937m and 984 are set into the proper state. This "early decode" makes sure that data addresses and codes will be forwarded to the

proper cross-links in all cases. Without such early decode logic, the cross-links could all be in a state with their drivers disabled. If one at the memory controllers were also disabled, then its cross-links would never receive addresses, data and control codes, effectively disabling all the I/O modules connected to that cross-link.

Prior to describing the driver control signals generated by decode logic 970, 971, 970m, 971m, and 998, it is necessary to understand the different modes that these zones, and therefore the cross-links 90 and 95, can be in. Fig. 13 contains a diagram of the different states A-F, and a table explaining the states which correspond to each mode.

At start-up and in other instances, both zones are in state A which is known as the OFF mode for both zones. In that mode, the computer systems in both zones are operating independently. After one of the zones' operating system requests the ability to communicate with the I/O of the other zone, and that request is honored, then the zones enter the master/slave mode, shown as states B and C. In such modes, the zone which is the master, has an operating CPU and has control of the I/O modules of its zone and of the other zone.

Upon initiation of resynchronization, the computer system leaves the master/slave modes, either states B or C, and enters a resync slave/resync master mode, which is shown as states E and F. In those modes, the zone that was the master zone is in charge of bringing the CPU of the other zone on line. If the resynchronization fails, the zones revert to the same master/slave mode that they were in prior to the resynchronization attempt.

If the resynchronization is successful, however, then the zones enter state D, which is the full duplex mode. In this mode, both zones are operating together in lockstep synchronization. Operation continues in this mode until there is a CPU/MEM fault, in which case the system enters one of the two master/slave modes. The slave is the zone whose processor experienced the CPU/MEM fault.

When operating in state D, the full duplex mode, certain errors, most notably clock phase errors, necessitate splitting the system into two independent processing systems. This causes system 10 to go back into state A.

Decode logic 970, 970m, 971, 971m, and 998 (collectively referred to as the cross-link control logic), which are shown in Figs. 11 and 12, have access to the resync mode bits 915 and the cross-link mode bits 916, which are shown in Fig. 10, in order to determine how to set the cross-link drivers and multiplexers into the proper states. In addition, the cross-link decode logic also receives and analyzes a portion of an address sent from memory controllers 70 and 75 during data transactions to extract addressing information that further indicates to the cross-link decode logic how to set the state of the cross-link multiplexers and drivers.

The information needed to set the states of the multiplexers is fairly straightforward once the different

modes and transactions are understood. The only determination to be made is the source of the data. Thus when cross-links 90 and 95 are in the slave mode, multiplexers 935, 935m, and 966 will select data addresses and codes from zone 11'. Those multiplexers will also select data, addresses and codes from the other zone if cross-links 90 and 95 are in full duplex mode, the address of an I/O instruction is for a device connected to an I/O module in zone 11, and the cross-link with the affected multiplexer is in a cross-over mode. In a cross-over mode, the data to be sent on the module interconnect is to be received from the other zone for checking. In the preferred embodiment, module interconnect 130 would receive data, addresses and codes from the primary rail in zone 11 and module interconnect would receive data, addresses and codes from the mirror rail in zone 11'. Alternatively, module interconnect 132 could receive data, addresses and codes from the primary rail in zone 11' which would allow the primary rail of one zone to be compared with the mirror rail of the other zone.

Multiplexers 945, 945m, and 982 will be set to accept data, address and codes from whichever zone is the source of the data. This is true both when all the cross-links are in full duplex mode and the data, address and codes are received from I/O modules and when the cross-link is in a resync slave mode and the data, address and codes are received from the memory controllers of the other zone.

If the addressing information from memory controllers 70 and 75 indicates that the source of response data and codes is the cross-link's own parallel registers 910, then multiplexers 942, 942m, and 974 are set to select data and codes from those registers. Similarly, if the addressing information from memory controllers 70 and 75 indicates that the source of response data is the cross-link's own serial register 920, then multiplexers 949 and 949m are set to select data and codes from those registers.

Multiplexers 949 and 949m are also set to select data from decode logic 970 and 970m, respectively, if the information is a control code during memory resync operations, and to select the ERR code if the EXCLUSIVE OR gates 960 and 960m identify a miscompare between the data transmitted via cross-links 90 and 95. In this latter case, the control of the multiplexers 949 and 949m is generated from the EXCLUSIVE OR gates 960 and 960m rather than from the cross-link control logic. Multiplexers 949 and 949m also select codes from serial cross-link registers 910 when those registers are requested or the output of multiplexers 945 and 945m when those codes are requested. Multiplexers 945 and 945m select either the outputs from multiplexers 942 and 942m, respectively, or I/O codes from cross-links 90' and 95', respectively.

Multiplexer 976 selects either data and addresses from module interconnect 130 in the case of a transaction with an I/O module, or data and addresses from

memory controller 90 when the data and addresses are to be sent to cross-link 90' either for I/O or during memory resynchronization.

Drivers 937 and 937m are activated when cross-

5 links 90 and 95 are in duplex, master or resync master modes. Drivers 940 and 940m are activated for I/O transactions in zone 11. Drivers 946 and 946m are activated when cross-links 90 and 95 are in the duplex or slave modes. Drivers 951 and 951m are always activated.

10 Driver 969 is activated during I/O writes to zone 11. Driver 984 is activated when cross-link 90 is sending data and addresses to I/O in zone 11', or when cross-link 90 is in the resync master mode. Receiver 986 receives 15 data from cross-link 90'. Drivers 992 and 994 are activated when data is being sent to memory controller 70; driver 994 is activated when the contents of the serial cross-link register 910 are read and driver 992 is activated during all other reads.

20 5. Oscillator

When both processing systems 20 and 20' are each 25 performing the same functions in the full duplex mode, it is imperative that CPU modules 30 and 30' perform operations at the same rate. Otherwise, massive amounts of processing time will be consumed in resynchronizing processing systems 20 and 20' for I/O and interprocessor error checking. In the preferred embodiment of processing systems 20 and 20', their basic clock signals are synchronized and phase-locked to each other. The fault tolerant computing system 10 includes a timing system to control the frequency of the clock signals to processing systems 20 and 20' and to minimize 30 the phase difference between the clock signals for each processing system.

35 Fig. 14 shows a block diagram of the timing system of this invention embedded in processing systems 20 and 20'. The timing system comprises oscillator system 200 in CPU module 30 of processing system 20, and oscillator system 200' in CPU module 30' of processing system 20'. The elements of oscillator 200' are equivalent to those for oscillator 200 and both oscillator systems' operation is the same. Thus, only the elements 40 and operation of oscillator system 200 will be described, except if the operations of oscillator systems 200 and 200' differ.

45 As Fig. 14 shows, much of oscillator system 200, specifically the digital logic, lies inside of cross-link 95, but that placement is not required for the present invention. Oscillator system 200 includes a voltage-controlled crystal oscillator (VCXO) 205 which generates a basic oscillator signal preferably at 66.66 Mhz. The frequency of VCXO 205 can be adjusted by the voltage level at the 50 input.

55 Clock distribution chip 210 divides down the basic oscillator signal and preferably produces four primary clocks all having the same frequency. For primary CPU

40 the clocks are PCLK L and PCLK H, which are logical inverses of each other. For mirror CPU 50, clock distribution chip 210 produces clock signals MCLK L and MCLK H, which are also logical inverses of each other. The timing and phase relationship of these clock signals are shown in Fig. 15. Preferably, frequency of clock signals PCLK L, PCLK H, MCLK L, and MCLK H is about 33.33 Mhz. Clock chip 210 also produces a phase-locked loop signal CLKC H at 16.66 Mhz, also shown in Fig. 15. This phase locked loop signal is sent to clock logic 220 which buffers that signal.

Clock logic buffer 220 sends the CLKC H signal to oscillator 200' for use in synchronization. Clock logic buffer 220' in oscillator 200' sends its own buffered phase-locked loop signal CLKC' H to phase detector 230 in oscillator 200. Phase detector 230 also receives the buffered phase locked loop signal CLKC H from clock logic 220 through delay element 225. Delay element 225 approximates the delay due to the cable run from clock logic buffer 220'.

Phase detector 230 compares its input phase locked loop signals and generates two outputs. One is a phase differences signal 235 which is sent through loop amplifier 240 to the voltage input of VCXO 205. Phase differences signal 235 will cause amplifier 240 to generate a signal to alter the frequency of VCXO 205 to compensate for phase differences.

The other output of phase detector 230 is a phase error signal 236 which indicates possible synchronization faults.

Fig. 16 is a detailed diagram of phase detector 230. Phase detector 230 includes a phase comparator 232 and a voltage comparator 234. Phase comparator 232 receives the clock signal from delay element 225 (CLKC H) and the phase lock loop clock signal from oscillator 200' (CLKC' H) and generates phase differences signal 235 as a voltage level representing the phase difference of those signals.

If processing system 20 were the "slave" for purposes of clock synchronization, switch 245 would be in the "SLAVE" position (i.e., closed) and the voltage level 235, after being amplified by loop amplifier 240, would control the frequency of VCXO 205. If both switches 245 and 245' are in the "master" position, processing systems 20 and 20' would not be phase-locked and would be running asynchronously (independently).

The voltage level of phase differences signal 235 is also an input to voltage comparator 234 as are two reference voltages,  $V_{ref1}$  and  $V_{ref2}$ , representing acceptable ranges of phase lead and lag. If the phase difference is within tolerance, the PHASE ERROR signal will not be activated. If the phase difference is out of tolerance, then the PHASE ERROR signal 236 will be activated and sent to cross-link 95 via clock decoder 220.

## 6. I/O Module

Fig. 17 shows a preferred embodiment of an I/O

module 100. The principles of operation I/O module 100 are applicable to the other I/O modules as well.

Fig. 18 shows the elements in the preferred embodiment of firewall 1000. Firewall 1000 includes a 16 bit

5 bus interface 1810 to module interconnect 130 and a 32 bit bus interface 1820 for connection to bus 1020 shown in Fig 17. Interfaces 1810 and 1820 are connected by an internal firewall bus 1815 which also interconnects with the other elements of firewall 1000. Preferably bus 10 1815 is a parallel bus either 16 or 32 bits wide.

I/O module 100 is connected to CPU module 30 by means of dual rail module interconnects 130 and 132. Each of the module interconnects is received by firewalls 1000 and 1010, respectively. One of the firewalls, 15 which is usually, but not always firewall 1000, writes the data from module interconnect 130 onto bus 1020. The other firewall, in this case firewall 1010, checks that data against its own copy received from module interconnect 132 using firewall comparison circuit 1840 shown in Fig. 20 18. That checking is effective due to the lockstep synchronization of CPU modules 30 and 30' which causes data written to I/O module 100 from CPU modules 30 and 30' to be available at firewalls 1000 and 1010 substantially simultaneously.

25 Firewall comparison circuit 1840 only checks data received from CPU modules 30 and 30'. Data sent to CPU modules 30 and 30' from an I/O device have a common origin and thus do not require checking. Instead, data received from an I/O device to be sent to CPU modules 30 and 30' is checked by an error detection code (EDC), such as a cyclical redundancy check (CRC), which is performed by EDC/CRC generator 1850. EDC/CRC generator 1850 is also coupled to internal firewall bus 1815.

30 35 EDC/CRC generator 1850 generates and checks the same EDC/CRC code that is used by the I/O device. Preferably, I/O module 100 generates two EDC. One, which can also be a EDC/CRC, is used for an interface to a network, such as the Ethernet packet network to 40 which module 100 is coupled (see element 1082 in Fig. 17). The other is used for a disk interface such as disk interface 1072 in Fig. 17.

45 EDC/CRC coverage is not required between CPU module 30 and I/O module 100 because the module interconnects are duplicated. For example in CPU module 30, cross-link 90 communicates with firewall 1000 through module interconnect 130, and cross-link 95 communicates with firewall 1010 through module interconnect 132.

50 55 A message received from Ethernet network 1082 is checked for a valid EDC/CRC by network control 1080 shown in Fig. 17. The data, complete with EDC/CRC, is written to a local RAM 1060 also shown in Fig. 17. All data in local RAM 1060 is transferred to memory module 60 using DMA. A DMA control 1890 coordinates the transfer and directs EDC/CRC generator 1850 to check the validity of the EDC/CRC encoded data being transferred.

Most data transfers with an I/O device are done with DMA. Data is moved between main memory and I/O buffer memory. When data is moved from the main memory to an I/O buffer memory, an EDC/CRC may be appended. When the data is moved from I/O buffer memory to main memory, an EDG/CRC may be checked and moved to main memory or may be stripped. When data is moved from the I/O buffer memory through an external device, such as a disk or Ethernet adaptor the EDC/CRC may be checked locally or at a distant receiving node, or both. The memory data packets may have their EDC/CRC generated at the distant node or by the local interface on the I/O module.

This operation ensures that data residing in or being transferred through a single rail system like I/O module 100 is covered by an error detection code, which is preferably at least as reliable as the communications media the data will eventually pass through. Different I/O modules, for example those which handle synchronous protocols, preferably have an EDC/CRC generator which generates and checks the EDC/CRC codes of the appropriate protocols.

In general, DMA control 1890 handles the portion of a DMA operation specific to the shared memory controller 1050 and local RAM 1060 being addressed. The 32 bit bus 1020 is driven in two different modes. During DMA setup, DMA control 1890 uses bus 1020 as a standard asynchronous microprocessor bus. The address in local RAM 1060 where the DMA operation will occur is supplied by shared memory controller 1050 and DMA control 1890. During the actual DMA transfer, DMA control 1890 directs DMA control lines 1895 to drive bus 1020 in a synchronous fashion. Shared memory controller 1050 will transfer a 32 bit data word with bus 1020 every bus cycle, and DMA control 1890 keeps track of how many words are left to be transferred. Shared memory control 1050 also controls local RAM 1060 and creates the next DMA address.

The I/O modules (100, 110, 120) are responsible for controlling the read/write operations to their own local RAM 1060. The CPU module 30 is responsible for controlling the transfer operations with memory array 60. The DMA engine 800 of memory controllers 70 and 75 (shown in Fig. 8) directs the DMA operations on the CPU module 30. This division of labor prevents a fault in the DMA logic on any module from degrading the data integrity on any other module in zones 11 or 11'.

The functions of trace RAM 1872 and trace RAM controller 1870 are described in greater detail below. Briefly, when a fault is detected and the CPUs 40, 40', 50 and 50' and CPU modules 30 and 30' are notified, various trace RAMs throughout computer system 10 are caused to perform certain functions described below. The communications with the trace RAMs takes place over trace bus 1095. Trace RAM control 1870, in response to signals from trace bus 1095, causes trace RAM 1872 either to stop storing, or to dump its contents over trace bus 1095.

I/O module bus 1020, which is preferably a 32 bit parallel bus, couples to firewalls 1000 and 1010 as well as to other elements of the I/O module 100. A shared memory controller 1050 is also coupled to I/O bus 1020

5 in I/O module 100. Shared memory controller 1050 is coupled to a local memory 1060 by a shared memory bus 1065, which preferably carries 32 bit data. Preferably, local memory 1060 is a RAM with 256 Kbytes of memory, but the size of RAM 1060 is discretionary. The shared memory controller 1050 and local RAM 1060 provide memory capability for I/O module 100.

Disk controller 1070 provides a standard interface to a disk, such as disks 1075 and 1075' in Fig. 1. Disk controller 1070 is also coupled to shared memory controller 1050 either for use of local RAM 1060 or for communication with I/O module bus 1020.

A network controller 1080 provides an interface to a standard network, such as the ETHERNET network, by way of network interface 1082. Network controller

20 1080 is also coupled to shared memory controller 1050 which acts as an interface both to local RAM 1060 and I/O module bus 1020. There is no requirement, however, for any one specific organization or structure of I/O module bus 1020.

25 PCIM (power and cooling interface module) support element 1030 is connected to I/O module bus 1020 and to an ASCII interface 1032. PCIM support element 1030 allows processing system 20 to monitor the status of the power system (i.e., batteries, regulators, etc.) and the cooling system (i.e., fans) to ensure their proper operation. Preferably, PCIM support element 1030 only receives messages when there is some fault or potential fault indication, such as an unacceptably low battery voltage. It is also possible to use PCIM support element 30 1030 to monitor all the power and cooling subsystems periodically. Alternatively PCIM support element 1030 may be connected directly to firewall S 1000 and 1010.

35 Diagnostics microprocessor 1100 is also connected to the I/O module bus 1020. In general, diagnostics microprocessor 1100 is used to gather error checking information from trace RAMS, such as trace RAM 1872, when faults are detected. That data is gathered into trace buses 1095 and 1096, through firewalls 1000 and 1010, respectively, through module bus 1020, and into 40 microprocessor 1100.

#### D. INTERPROCESSOR AND INTERMODULE COMMUNICATION

##### 50 1. Data Paths

The elements of computer system 10 do not by themselves constitute a fault tolerant system. There needs to be a communications pathway and protocol 55 which allows communication during normal operations and operation during fault detection and correction. Key to such communication is cross-link pathway 25. Cross-link pathway 25 comprises the parallel links, serial links,

and clock signals already described. These are shown in Fig. 19. The parallel link includes two identical sets of data and address lines, control lines, interrupt lines, coded error lines, and a soft reset request line. The data and address lines and the control lines contain information to be exchanged between the CPU modules, such as from the module interconnects 130 and 132 (or 130' and 132') or from memory module 60 (60').

The interrupt lines preferably contain one line for each of the interrupt levels available to I/O subsystem (modules 100, 110, 120, 100', 110' and 120'). These lines are shared by cross-links 90, 95, 90' and 95'.

The coded error lines preferably include codes for synchronizing a console "HALT" request for both zones, one for synchronizing a CPU error for both zones, one for indicating the occurrence of a CPU/memory failure to the other zone, one for synchronizing DMA error for both zones, and one for indicating clock phase error. The error lines from each zone 11 or 11' are inputs to an OR gate, such as OR gate 1990 for zone 11 or OR gate 1990' for zone 11'. The output at each OR gate provides an input to the cross-links of the other zone.

The fault tolerant processing system 10 is designed to continue operating as a dual rail system despite transient faults. The I/O subsystem (modules 100, 110, 120, 100', 110', 120') can also experience transient errors or faults and continue to operate. In the preferred embodiment, an error detected by firewall comparison circuit 1840 will cause a synchronized error report to be made through pathway 25 for CPU directed operations. Hardware in CPU 30 and 30' will cause a synchronized soft reset through pathway 25 and will retry the faulted operation. For DMA directed operations, the same error detection results in synchronous interrupts through pathway 25, and software in CPUs 40, 50, 40' and 50' will restart the DMA operation.

Certain transient errors are not immediately recoverable to allow continued operation in a full-duplex, synchronized fashion. For example, a control error in memory module 60 can result in unknown data in memory module 60. In this situation, the CPUs and memory elements can no longer function reliably as part of a fail safe system so they are removed. Memory array 60 must then undergo a memory resync before the CPUs and memory elements can rejoin the system. The CPU/memory fault code of the coded error lines in pathway 25 indicates to CPU 30' that the CPUs and memory elements of CPU 30 have been faulted.

The control lines, which represent a combination of cycle type, error type, and ready conditions, provide the handshaking between CPU modules (30 and 30') and the I/O modules. Cycle type, as explained above, defines the type of bus operation being performed: CPU I/O read, DMA transfer, DMA setup, or interrupt vector request. Error type defines either a firewall miscompare or a CRC error. "Ready" messages are sent between the CPU and I/O modules to indicate the completion of requested operations.

The serial cross-link includes two sets of two lines to provide a serial data transfer for a status read, loop-back, and data transfer.

5 The clock signals exchanged are the phase locked clock signals CLKC H and CLKC' H (delayed).

Figs. 20A-D show block diagrams of the elements of CPU modules 30 and 30' and I/O modules 100 and 100' through which data passes during the different operations. Each of those elements has each been described previously.

Fig. 20A shows the data pathways for a typical CPU I/O read operation of data from an I/O module 100, such as a CPU I/O register read operation of register data from shared memory controller 1050 (1050'). Such an operation will be referred to as a read of local data, to distinguish it from a DMA read of data from local memory 1060, which usually contains data from an internal device controller. The local data are presumed to be stored in local RAM 1060 (1060') for transfer through shared memory controller 1050 (1050'). For one path, the data pass through firewall 1000, module interconnect 130, to cross-link 90. As seen in Fig. 12, cross-link 90 delays the data from firewall 1000 to memory controller 70 so that the data to cross-link 90' may be presented to memory controller 70 at the same time the data are presented to memory controller 70, thus allowing processing systems 20 and 20' to remain synchronized. The data then proceed out of memory controllers 70 and 70' into CPUs 40 and 40' by way of internal busses 46 and 46'.

30 A similar path is taken for reading data into CPUs 50 and 50'. Data from the shared memory controller 1050 proceeds through firewall 1010 and into cross-link 95. At that time, the data are routed both to cross-link 95' and through a delay unit inside cross-link 95.

35 CPU I/O read operations may also be performed for data received from the I/O devices of processing system 20' via a shared memory controller 1050' and local RAM in I/O device 100'.

Although I/O modules 100, 110, and 120 are similar 40 and correspond to I/O modules 100', 110', and 120', respectively, the corresponding I/O modules are not in lockstep synchronization. Using memory controller 1050' and local RAM 1060' for CPU I/O read, the data would first go to cross-links 90' and 95'. The remaining data path is equivalent to the path from memory controller 1050. The data travel from the cross-links 90' and 95' up through memory controllers 70' and 75' and finally to CPUs 40' and 50', respectively. Simultaneously, the data travel across to cross-links 90 and 95, respectively, and then, without passing through a delay element, the data continue up to CPUs 40 and 50, respectively.

55 Fig. 20B shows a CPU I/O write operation of local data. Such local data are transferred from the CPUs 40, 50, 40' and 50' to an I/O module, such as I/O module 100. An example of such an operation is a write to a register in shared memory controllers 1050. The data transferred by CPU 40 proceed along the same path but

in a direction opposite to that of the data during the CPU I/O read. Specifically, such data pass through bus 46, memory controller 70, various latches (to permit synchronization), firewall 1000, and memory controller 1050. Data from CPU 50' also follow the path of the CPU I/O reads in a reverse direction. Specifically, such data pass through bus 56', memory controller 75', cross-link 95', cross-link 95, and into firewall 1010. As indicated above, firewalls 1000 and 1010 check the data during I/O write operations to check for errors prior to storage.

When writes are performed to an I/O module in the other zone, a similar operation is performed. However, the data from CPUs 50 and 40' are used instead of CPUs 50' and 40.

The data from CPUs 50 and 40' are transmitted through symmetrical paths to shared memory controller 1050'. The data from CPUs 50 and 40' are compared by firewalls 1000' and 1010'. The reason different CPU pairs are used to service I/O write data is to allow checking of all data paths during normal use in a full duplex system. Interrail checks for each zone were previously performed at memory controllers 70, 75, 70' and 75'.

Fig. 20C shows the data paths for DMA read operations. The data from memory array 600 pass simultaneously into memory controllers 70 and 75 and then to cross-links 90 and 95. Cross-link 90 delays the data transmitted to firewall 1000 so that the data from cross-links 90 and 95 reach firewalls 1000 and 1010 at substantially the same time.

Similar to the CPU I/O write operation, there are four copies of data to the various cross-links. At the firewall, only two copies are received. A different pair of data are used when performing reads to zone 11. The data paths for the DMA write operation are shown in Fig. 20D and are similar to those for a CPU I/O read. Specifically, data from shared memory controller 1050' proceed through firewall 1000', cross-link 90' (with a delay), memory controller 70', and into memory array 600'. Simultaneously, the data pass through firewall 1010', cross-link 95' (with a delay), and memory controller 75', at which time it is compared with the data from memory controller 70' during an interrail error check. As with the CPU I/O read, the data in a DMA write operation may alternatively be brought up through shared memory controller 1050 in an equivalent operation.

The data out of cross-link 90' also pass through cross-link 90 and memory controller 70 and into memory array 600. The data from cross-link 95' pass through cross-link 95 and memory controller 75, at which time they are compared with the data from memory controller 70' during a simultaneous interrail check.

The data path for a memory resynchronization (resync) operation is shown in Fig. 20E. In this operation the contents of both memory arrays 60 and 60' must be set equal to each other. In memory resync, data from memory array 600' pass through memory controllers 70' and 75' under DMA control, then through cross-links 90' and 95', respectively. The data then enters cross-links

90 and 95 and memory controllers 70 and 75, respectively, before being stored in memory array 600.

## 2. Resets

The preceding discussions of system 10 have made reference to many different needs for resets. In certain instances not discussed, resets are used for standard functions, such as when power is initially applied to system 10. Most systems have a single reset which always sets the processor back to some predetermined or initial state, and thus disrupts the processors' instruction flow. Unlike most other systems, however, resets in system 10 do not affect the flow of instruction execution by CPUs 40, 40', 50 and 50' unless absolutely necessary. In addition, resets in system 10 affect only those portions that need to be reset to restore normal operation.

Another aspect of the resets in system 10 is their containment. One of the prime considerations in a fault tolerant system is that no function should be allowed to stop the system from operating should that function fail. For this reason, no single reset in system 10 controls elements of both zones 11 and 11' without direct cooperation between zones 11 and 11'. Thus, in full duplex mode of operation, all resets in zone 11 will be independent of resets in zone 11'. When system 10 is in master/slave mode, however, the slave zone uses the resets of the master zone. In addition, no reset in system 10 affects the contents of memory chips. Thus neither cache memory 42 and 52, scratch pad memory 45 and 55 nor memory module 60 lose any data due to a reset.

There are preferably three classes of resets in system 10: "clock reset," "hard reset," and "soft reset." A clock reset realigns all the clock phase generators in a zone. A clock reset in zone 11 will also initialize CPUs 40 and 50 and memory module 60. A clock reset does not affect the module interconnects 130 and 132 except to realign the clock phase generators on those modules. Even when system 10 is in master/slave mode, a clock reset in the slave zone will not disturb data transfers from the master zone to the slave zone module interconnect. A clock reset in zone 11', however, will initialize the corresponding elements in zone 11'.

In general, a hard reset returns all state devices and registers to some predetermined or initial state. A soft reset only returns state engines and temporary storage registers to their predetermined or initial state. The state engine in a module is the circuitry that defines the state of that module. Registers containing error information and configuration data will not be affected by a soft reset. Additionally, system 10 will selectively apply both hard resets and soft resets at the same time to reset only those elements that need to be reinitialized in order to continue processing.

The hard resets clear system 10 and, as in conventional systems, return system 10 to a known configuration. Hard resets are used after power is applied, when zones are to be synchronized, or to initialize or disable

an I/O module. In system 10 there are preferably four hard resets: "power up reset," "CPU hard reset," "module reset," and "device reset." Hard resets can be further broken down into local and system hard resets. A local hard reset only affects logic that responds when the CPU is in the slave mode. A system hard reset is limited to the logic that is connected to cross-link cables 25 and module interconnects 130 and 132.

The power up reset is used to initialize zones 11 and 11' immediately after power is supplied. The power up reset forces an automatic reset to all parts of the zone. A power up reset is never connected between the zones of system 11 because each zone has its own power supply and will thus experience different length "power-on" events. The power up reset is implemented by applying all hard resets and a clock reset to zone 11 or 11'.

The CPU hard reset is used for diagnostic purposes in order to return a CPU module to a known state. The CPU hard reset clears all information in the CPUs, memory controllers, and memory module status registers in the affected zone. Although the cache memories and memory modules are disabled, the contents of the scratch pad RAMs 45 and 55 and of the memory module 60 are not changed. In addition, unlike the power up reset, the CPU hard reset does not modify the zone identification of the cross-links nor the clock mastership. The CPU hard reset is the sum of all local hard resets that can be applied to a CPU module and a clock reset.

The module hard reset is used to set the I/O modules to known state, such as during bootstrapping, and is also used to remove a faulting I/O module from the system. The I/O module hard reset clears everything on the I/O module, leaves the firewalls in a diagnostic mode, and disables the drivers.

A device reset is used to reset I/O devices connected to the I/O modules. The resets are device dependent and are provided by the I/O module to which the device is connected.

The other class of resets is soft resets. As explained above, soft resets clear the state engines and temporary registers in system 10 but they do not change configuration information, such as the mode bits in the cross-links. In addition, soft resets also clear the error handling mechanisms in the modules, but they do not change error registers such as system error register 898 and system fault address register 865.

Soft resets are targeted so that only the necessary portions of the system are reset. For example, if module interconnect 130 needs to be reset, CPU 40 is not reset nor are the devices connected to I/O module 100.

There are three unique aspects of soft resets. One is that each zone is responsible for generating its own reset. Faulty error or reset logic in one zone is thus prevented from causing resets in the non-faulted zone.

The second aspect is that the soft reset does not disrupt the sequence of instruction execution. CPUs 40, 40', 50, 50' are reset on a combined clock and hard reset only. Additionally memory controllers 70, 75, 70' and 75'

have those state engines and registers necessary to service CPU instructions attached to hard reset. Thus the soft reset is transparent to software execution.

The third aspect is that the range of a soft reset, that

- 5 is the number of elements in system 10 that is affected by a soft reset, is dependent upon the mode of system 10 and the original reset request. In full duplex mode, the soft reset request originating in CPU module 30 will issue a soft reset to all elements of CPU module 30 as well as all firewalls 1000 and 1010 attached to module interconnect 130 and 132. Thus all modules serviced by module interconnect 130 and 132 will have their state engines and temporary registers reset. This will clear the system pipeline of any problem caused by a transient error. Since system 10 is in duplex mode, zone 11' will be doing everything that zone 11 is. Thus CPU module 30' will, at the same time as CPU module 30, issue a soft reset request. The soft reset in zone 11' will have the same effect as the soft reset in zone 11.
- 10
- 15
- 20 When system 10 is in a master/slave mode, however, with CPU module 30' in the slave mode, a soft reset request originating in CPU module 30 will, as expected, issue a soft reset to all elements of CPU module 30 as well as all firewalls 1000 and 1010 attached to module interconnects 130 and 132. Additionally, the soft reset request will be forwarded to CPU module 30' via cross-links 90 and 90', cross-link cables 25, and cross-links 90' and 95'. Parts of module interconnects 130' and 132' will receive the soft reset. In this same configuration, a soft reset request originating from CPU module 30' will only reset memory controllers 70' and 75' and portions of cross-links 90' and 95'.
- 25
- 30

Soft resets include "CPU soft resets" and "system soft resets."

- 35 A CPU soft reset is a soft reset that affects the state engines on the CPU module that originated the request. A system soft reset is a soft reset over the module interconnect and those elements directly attached to it. A CPU module can always request a CPU soft reset. A system soft reset can only be requested if the cross-link of the requesting CPU is in duplex mode, master/slave mode, or off mode. A cross-link in the slave mode will take a system soft reset from the other zone and generate a system soft reset to its own module interconnects.
- 40
- 45

CPU soft resets clear the CPU pipeline following an error condition. The CPU pipeline includes memory interconnects 80 and 82, latches (not shown) in memory controllers 70 and 75, DMA engine 800, and cross-links 90 and 95. The CPU soft reset can also occur following a DMA or I/O time-out. A DMA or I/O time-out occurs when the I/O device does not respond within a specified time period to a DMA or an I/O request.

- 50 Fig. 21 shows the reset lines from the CPU modules 30 and 30' to the I/O modules 100, 110, 100', and 110' and to the memory modules 60 and 60'. The CPU module 30 receives a DC OK signal indicating when the power supply has settled. It is this signal which initializes the power-up reset. CPU module 30' receives a similar sig-
- 55

nal from its power supply.

One system hard reset line is sent to each I/O module, and one system soft reset is sent to every three I/O modules. The reason that single hard reset is needed for each module is because the system hard reset line are used to remove individual I/O modules from system 10. The limitation of three I/O modules for each system soft reset is merely a loading consideration. In addition, one clock reset line is sent for every I/O module and memory module. The reason for using a single line per module is to control the skew by controlling the load.

Fig. 22 shows the elements of CPU module 30 which relate to resets. CPUs 40 and 50 contain clock generators 2210 and 2211, respectively. Memory controllers 70 and 75 contain clock generators 2220 and 2221, respectively, and cross-links 90 and 95 contain clock generators 2260 and 2261, respectively. The clock generators divide down the system clock signals for use by the individual modules.

Memory controller 70 contains reset control circuitry 2230 and a soft reset request register 2235. Memory controller 75 contains reset control circuitry 2231 and a soft reset request register 2236.

Cross-link 90 contains both a local reset generator 2240 and a system reset generator 2250. Cross-link 95 contains a local reset generator 2241 and a system reset generator 2251. The "local" portion of a cross-link is that portion of the cross-link which remains with the CPU module when that cross-link is in the slave mode and therefore includes the serial registers and some of the parallel registers. The "system" portion of a cross-link is that portion of the cross-link that is needed for access to module interconnects 130 and 132 (or 130' and 132') and cross-link cables 25.

The local reset generators 2240 and 2241 generate resets for CPU module 30 by sending hard and soft reset signals to the local reset control circuits 2245 and 2246 of cross-links 90 and 95, respectively, and to the reset control circuits 2230 and 2231 of memory controller 70 and 75, respectively. Local cross-link reset control circuits 2245 and 2246 respond to the soft reset signals by resetting their state engines, the latches storing data to be transferred, and their error registers. Those circuits respond to the hard reset signals by taking the same actions as are taken for the soft resets, and by also resetting the error registers and the configuration registers. Reset control circuits 2230 and 2231 respond to hard and soft reset signals in a similar manner.

In addition, the local reset generator 2240 sends clock reset signals to the I/O modules 100, 130 and 120 via module interconnects 130 and 132. The I/O modules 100, 110, and 120 use the clock reset signals to reset their clocks in the manner described below. Soft reset request registers 2235 and 2236 send soft request signals to local reset generators 2240 and 2241, respectively.

System reset generators 2250 and 2251 of cross-links 90 and 95, respectively, send system hard reset

signals and system soft reset signals to I/O modules 100, 130, and 120 via module interconnects 130 and 132, respectively. I/O modules 100, 110, and 120 respond to the soft reset signals by resetting all registers that are dependent on CPU data or commands. Those modules respond to the hard reset signals by resetting the same register as soft resets do, and by also resetting any configuration registers.

In addition, the system reset generators 2250 and 2251 also send the system soft and system hard reset signals to the system reset control circuit 2255 and 2256 of each cross-link. System reset control circuit 2255 and 2256 respond to the system soft reset signals and to the system hard reset signals in a manner similar to the response of the local reset control circuits to the local soft and local hard reset signals.

Memory controllers 70 and 75 cause cross-links 90 and 95, respectively, to generate the soft resets when CPUs 40 and 50, respectively, write the appropriate codes into soft reset request registers 2235 and 2236, respectively. Soft reset request registers 2235 and 2236 send soft reset request signals to local reset generators 2240 and 2241, respectively. The coded error signal is sent from memory controller 70 to local reset generators 2240 and 2241.

System soft resets are sent between zones along the same data paths data and control signals are sent. Thus, the same philosophy of equalizing delays is used for resets as for data and addresses, and resets reach all of the elements in both zones at approximately the same time.

Hard resets are generated by CPUs 40 and 50 writing the appropriate code into the local hard reset registers 2243 or by the request for a power up reset caused by the DC OK signal.

Synchronization circuit 2270 in cross-link 90 includes appropriate delay elements to ensure that the DC OK signal goes to all of the local and reset generators 2240, 2250, 2241 and 2251 at the same time.

In fact, synchronization of resets is very important in system 10. That is why the reset signals originate in the cross-links. In that way, the resets can be sent to arrive at different modules and elements in the modules approximately synchronously.

With the understanding of the structure in Figs. 21 and 22, the execution of the different hard resets can be better understood. The power up reset generates both a system hard reset, a local hard reset and a clock reset. Generally, cross-links 90, 95, 90' and 95' are initially in both the cross-link off and resync off modes, and with both zones asserting clock mastership.

The CPU/MEM fault reset is automatically activated whenever memory controllers 70, 75, 70' and 75' detect a CPU/MEM fault. The coded error logic is sent from error logic 2237 and 2238 to both cross-links 90 and 95. The CPU module generated the fault is then removed from system 10 by setting its cross-link to the slave state and by setting the cross-link in the other CPU module

to the master state. The non-faulting CPU module will not experience a reset, however. Instead, it will be notified of the fault in the other module through a code in a serial cross-link error register (not shown). The CPU/MEM fault reset consists of a clock reset to the zone with the failing CPU module and a local soft reset to that module.

A resync reset is essentially a system soft reset with a local hard reset and a clock reset. The resync reset is used to bring two zones into lockstep synchronization. If, after a period in which zones 11 and 11' were not synchronized, the contents of the memory modules 60 and 60', including the stored states of the CPU registers, are set equal to each other, the resync reset is used to bring the zones into a compatible configuration so they can restart in a duplex mode.

The resync reset is essentially a CPU hard reset and a clock reset. The resync reset is activated by software writing the resync reset address into one of the parallel cross-link registers. At that time, one zone should be in the cross-link master/resync master mode and the other in the cross-link slave/resync slave mode. A simultaneous reset will then be performed on both the zones which, among other things, will set all four cross-links into the duplex mode. Since the resync reset is not a system soft reset, the I/O modules do not receive reset.

The preferred embodiment of system 10 also ensures that clock reset signals do not reset conforming clocks, only non-conforming clocks. The reason for this is that whenever a clock is reset, it alters the timing of the clocks which in turn affects the operation of the modules with such clocks. If the module was performing correctly and its clock was in the proper phase, then altering its operation would be both unnecessary and wasteful.

Fig. 23 shows a preferred embodiment of circuitry which will ensure that only nonconforming clocks are reset. The circuitry shown in Fig. 23 preferably resides in the clock generators 2210, 2211, 2220, 2221, 2260, and 2261 of the corresponding modules shown in Fig. 22.

In the preferred embodiment, the different clock generators 2210, 2211, 2220, 2221, 2260, and 2261 include a rising edge detector 2300 and a phase generator 2310. The rising edge detector 2300 receives the clock reset signals from the cross-links 90 and 95 and generates a pulse of known duration concurrent with the rising edge of the clock reset signal. That pulse is an input to the phase generator 2310 as are the internal clock signals for the particular module. The internal clock signals for that module are clock signals which are derived from the system clock signals that have been distributed from oscillator systems 200 and 200'.

Phase generator 2310 is preferably a divide-down circuit which forms different phases for the clock signals. Other designs for phase generator 2310, such as recirculating shift registers, can also be used.

Preferably, the rising edge pulse from rising edge detector 2300 causes phase generator 2310 to output

a preselected phase. Thus, for example, if phase generator 2310 were a divide-down circuit with several stages, the clock reset rising edge pulse could be a set input to the stage which generates the preselected phase and

5 a reset input to all other stages. If phase generator 2310 were already generating that phase, then the presence of the synchronized clock reset signal would be essentially transparent.

The resets thus organized are designed to provide

10 the minimal disruption to the normal execution of system 10, and only cause the drastic action of interrupting the normal sequences of instruction execution when such drastic action is required. This is particularly important in a dual or multiple zone environment because of the  
15 problems of resynchronization which conventional resets cause. Thus, it is preferable to minimize the number of hard resets, as is done in system 10.

#### E. RESYNCHRONIZATION AND BULK MEMORY 20 TRANSFER

If zones 11 and 11' are not in lockstep synchronization and such synchronization is desired for full duplex operation, the first step that needs to occur is to bring  
25 the memory modules 60 and 60' of zones 11 and 11', respectively, to the same state. This involves ensuring that the contents of memory arrays 600 and 600' are identical. The data flow for a bulk transfer of contents from array 600' to array 600 is shown in Fig. 20E.

30 In accordance with this invention, transfer of memory contents between zones occurs without stopping zone 11 (assumed to be the operating zone) and with minimal interruption of instruction execution. Thus applications running in zone 11 can continue running during the bulk memory transfer. A preferred method in accordance with the invention is shown in Fig. 24.

Before entering the bulk memory transfer and resynchronization operation 2400, however, CPUs 40 and 50 must have audited the pages of the master memory array 600 and slave memory array 600' to determine  
40 which pages were present. CPUs 40 and 50 would likely have already completed an audit of memory array 600 in the course of routine memory management operations. CPUs 40 and 50 would need to perform an audit  
45 of array 600' unless the audit information had somehow been previously given to CPUs 40 and 50. The bulk transfer operation would only take place if slave memory array 600' had a page for every page of master memory array 600.

50 As Fig. 24 shows, the first step in the bulk memory copy operation 2400 involves setting cross-links 90, 95, 90' and 95' to the memory resync master/slave mode to permit the data path shown in Fig. 20E (step 2410). This is done through software control of mode bits 915.

55 The effect on the control paths of cross-links 90' and 95' of being in the resync slave mode can be seen by referring to Fig. 11. The control paths define how the memory commands move from memory controllers 70

and 75 to memory controllers 70' and 75'. The primed numbers in the following discussion refers to elements of zone 11'. Those which are shown are equivalent to the unprimed elements which are shown. For bulk memory transfer, a multiplexer 935' and 935m' would be set to select the inputs from latch 938' and 938m', respectively, and forward commands to decoders 970' and 970m'. If the commands were intended for the slave memory controllers, in this example, memory controllers 70' and 75', decoders 970' and 970m' modify the command from memory controllers 70 and 75. In the preferred embodiment, this is done by changing the code from a four-bit down code to a three-bit up code. Multiplexers 949' and 949m' would then select the outputs from decode logic 970 and 970m', respectively, and transmit those codes to memory controllers 70' and 75' through drivers 951' and 951m', respectively.

The path for data through cross-link 90' when in the cross-link slave mode can be seen by referring to Fig. 12. The data path through cross-link 95' uses corresponding elements. Data from cross-link cable 25 enters through receiver 986' into latch 988'. Multiplexer 982' selects the data from latch 988', and latch 990' stores that data. Driver 992' sends that stored data to memory controller 90'. As explained below, this configuration also allows DMA data from I/O devices to be stored into memory modules 60 and 60' simultaneously.

Next, the refresh circuitry in sequencer 617' of memory module 60' is disabled (step 2420). In the preferred embodiment, CPUs 40' and 50' send disable signals to sequencer 617' when zone 11' enters the slave mode. The reason for disabling refresh signals from memory module 60' to memory array 600' is to prevent a logjam from occurring if memory array 600' was executing a refresh operation while receiving data from zone 11. During a refresh operation, a memory array cannot store new data, so the "logjam" would persist until memory array 600' finished its refresh operation.

Following the disabling of refresh circuitry, the pages of memory from master memory module 60 are queued sequentially into DMA engine 800 of the memory controller (i.e., controllers 70 and 75) of the master zone 11 (step 2430). In the preferred embodiment, each DMA transfer includes an indication of the type of transfer (in this example, a READ) the starting address (in this example, a boundary address for the next page), the number of elements to be transferred (in this example, the page size), and an indication of the destination (in this example, slave zone memory). This data would be sent to DMA engine 800.

Because the bulk memory transfer method of this invention allows the storage of DMA data from I/O devices, any requests to store such data, which will be referred to as non-resync DMA transfers, are intermixed into the DMA queues. This is also what allows the applications currently running on system 10 to continue during the bulk memory transfer process. Queue entries for DMA data transfers to I/O modules 100, 110, 120,

100', 110' and 120' are unmodified from the non-resync mode of operation. Queue entries for DMA data from I/O modules to main memory 60 require the additional information of the boundary address and element size for the slave memory.

The cross-links will automatically route all DMA data going from I/O modules 100, 110, 120, 100', 110', and 120' to both the master and slave memory controllers 70, 75, 70' and 75'. The intermixing of application DMAs is interactive. As an application's need for DMA activity arises, its requests can be added into the pending queue (step 2440).

DMA engines 800 and 800' examine the next entry in the DMA queue (step 2450). If the entry is for a bulk memory transfer, the master zone memory controllers 70 and 75 are given the configuration information of starting address and the number of elements to transfer as part of a read operation (step 2451). The slave zone memory controllers 70' and 75' get the same information as part of a write operation (step 2452).

For I/O to memory transfers, master zone memory controllers 70 and 75 are given a destination address and a transfer count (step 2453), as are slave memory controllers 70' and 75' (step 2454). The I/O module is sent the starting address and transfer count (step 2455). For the memory controllers, this is a write operation. For the I/O module this is a read operation.

For a memory to I/O transfer, master memory controllers and 75 are sent the starting address and transfer count for a read operation (step 2456). The I/O module receives the destination address and transfer count for a write operation (2457). The slave memory controllers have no involvement in this type of DMA transfer.

After the DMA transfer is set up, that operation is executed (step 2460).

During the DMA transfer of memory contents from the master zone memory to the slave zone memory, or for that matter, during any of the DMA transfers, any storage commands which are received by the master zone memory controller are also transferred to the slave zone memory controller (step 2465). Such storage commands can include, for example, write commands caused by the currently running application which are received by the master zone memory controller from the master zone CPU. The storage commands also include memory refresh address commands sent to memory module 60 from memory controllers 70 and 75.

The transfer of refresh commands occurs because sequencer 617' in slave zone 11' is not generating refresh signals. Slave zone memory module 60' will thus be able to store the data received from the master zone memory module 60 without interruption due to refresh signals. The operation of slave memory module 60' will not be compromised, however, because it will receive refresh signals from master zone memory controller 70. Also, because the master zone sends its refresh signals to the slave zone, the slave and master zone memory arrays 600' and 600 will refresh at the same time.

At the conclusion of each page transfer, a determination is made whether all the pages from the master zone memory module 60 have been transferred (step 2470). If not, after a predetermined delay (step 2480), the next page in the queue is transferred. The delay is set by CPUs 40 and 50 by writing a code to a DMA ratio portion 877 of memory control status register 876 in master zone memory controllers 70 and 75 and the delay is implemented by controlling when CPUs 40 and 50 instruct DMA engine 800 to send storage commands. In the preferred embodiment the delay sets the maximum ratio of DMA activity between 1/8 and 8/8 of the memory bandwidth. During such delay, however, storage commands are still being sent to zone 11'.

The purpose of the delay, which in the preferred embodiment is adjustable, is to prevent the DMA transfer from slowing down the normal processing of the CPUs (i.e., CPUs 40 and 50) too much during the memory resync operation. If the processing envisioned during resynchronization is light, the delay can be set to a small value. If the processing envisioned during resynchronization is heavy, then the delay can be set to a large value.

When all the master zone memory modules have been transferred, the master zone CPUs (i.e., CPUs 40 and 50) store their CPU registers into memory array 600 (step 2485) for later rebooting operation. The storage of the CPU registers into the memory array 600 also stores them into the slave memory array 600' because the zones are still configured for resynchronization.

Next, both zones are reset (step 2490) using a "resync reset." The resync reset flushes the cache memories 42, 52, 42' and 52', places the zones 11 and 11' into the same initial state, and also brings the zones out of the resync state (i.e., cancels the resync mode) by resetting the cross-links into duplex mode. The resync reset thus automatically cancels the slave zone's mimicking of the master zone's storage commands, and re-enables the refresh circuitry in sequencer 617'. After the resync reset, the system 10 is booted (step 2495) and can be started in the synchronized full duplex mode.

## V. CONCLUSION

The present invention, as shown by its preferred embodiment, thus achieves significant advantages, such as minimizing resynchronization problems and delays for dual zone systems. It will be apparent to those skilled in the art that modifications and variations can be made in the methods and apparatus of this invention. The invention in its broader aspects is, therefore, not limited to the specific details, representative methods and apparatus, and illustrative examples shown and described. Accordingly, departure may be made from such details without departing from the scope of the invention as defined in the accompanying claims.

## Claims

1. A method of operating a dual zone computer system (10) having a master zone (11) and a slave zone (11') capable of operating in synchronism and communicating with each other, wherein said master and slave zones each include:

a central processing unit (CPU) (30,30'); a memory system (60,60'), wherein each of said memory systems responds to storage commands and access commands received from the CPU in the same zone and from sources other than the CPU in the same zone, the storage commands including commands to write data into said memory systems and refresh commands to refresh the memory systems;

a plurality of CPU registers which define the state of the corresponding CPU; and said method synchronizing the slave zone to the master zone while allowing continued operation of the master zone by the sequential steps of:

(a) transferring (2460) the contents of the master zone memory system (60) to the slave zone memory system (60');

(b) transferring (2465) to the slave zone memory system (60') all of the storage commands, including write commands and refresh commands, received by the master zone memory system during the transfer of the master zone memory system contents to the slave zone memory system;

(c) storing (2485) the contents of the master zone CPU registers into the master zone memory system and into the slave zone memory system;

(d) setting the memory systems and CPU's of the master zone and slave zone to the same initial state by resetting the zones followed by the substantially simultaneous steps of:

(i) transferring (2495), from the master zone memory system to the master zone CPU registers, the contents of the master zone CPU registers stored in the master zone memory system, and

(ii) transferring (2495), from the slave zone memory system to the slave zone CPU registers, the contents of the slave zone CPU registers stored in the slave zone memory system.

2. The method of claim 1 wherein:

said master zone memory system includes a

master zone memory controller (70,75) which receives all of the write commands to the master zone memory system and which generates refresh commands for said master zone memory system, and  
 5  
 said slave zone memory system includes a slave zone memory controller (70',75') which receives all of the write commands to the slave zone memory system and which generates refresh commands for said slave zone memory system,  
 10  
 wherein said method also includes the step of: preventing (2420) the slave zone memory controller from generating refresh commands during the transfer of the master zone memory system contents to the slave zone memory system.

3. The method of claim 1 wherein both said master zone memory system (60) and said slave zone memory system (60') have the capability for direct memory access (DMA) transfer, and wherein step (a) includes the substeps of:  
 15  
 (i) outputting the contents of the master zone memory system via DMA transfer, and  
 (ii) inputting the contents of the master zone memory system into the slave zone memory system via DMA transfer.  
 20

4. The method of claim 3 wherein the master zone (11) and slave zone (11') each including a cross-link circuit (90,95,90',95') for transferring information between said master zone and said slave zone, and wherein step (a) further includes the substeps of:  
 25  
 configuring (2410) the cross-link circuits of the master zone and the slave zone to allow for DMA transfer of the contents of the master zone memory system to the slave zone memory system and to transfer all of the storage commands received by said master zone memory system to said slave zone memory system, while the contents of the master zone memory system are being transferred to the slave zone memory system.  
 30

5. The method of claim 4 wherein step (d) further includes the substeps of:  
 35  
 reconfiguring the master zone and slave zone cross-link means to stop the transfer of all of the storage commands received by the master zone memory system to the slave zone memory system.  
 40

6. The method of claim 3 wherein step (a)(i) further includes the substeps of  
 45  
 outputting pages of said master zone memory, and waiting (2480) a predetermined period of time after the output of each page of said master zone memory.  
 50

7. A dual zone computer system (10) having a master zone (11) and a slave zone (11') capable of operating in synchronism and communicating with each other, and capable of operating in a synchronizing mode during which the computer system is attempting to enter said synchronized mode, wherein said master and slave zones each include:  
 55  
 a central processing unit (CPU) (30,30');  
 a memory system (60,60'), adapted to respond to storage commands and access commands received from the CPU in the same zone and from sources other than the CPU in the same zone, the storage commands including commands to write data into said memory systems and refresh commands to refresh the memory systems;  
 a plurality of CPU registers which define the state of the corresponding CPU;  
 said computer system further comprising means for synchronizing the slave zone to the master zone while allowing continued operation of the master zone comprising:  
 slave zone cross-link means (90',95'), coupled to the slave zone memory system (60'), for receiving storage commands and data from the master zone when said computer system is in said synchronizing mode;  
 master zone cross-link means (90,95) for routing data stored in said master zone memory system, storage commands received by said master zone memory system, and master zone CPU register data, to said slave zone when said computer system is in said synchronizing mode;  
 means, in said master zone memory system, for transferring data including master zone CPU register data from the master zone memory system to said plurality of master zone CPU registers;  
 means, in said slave zone, for transferring said received master zone CPU register data into a said corresponding plurality of slave zone CPU registers, substantially simultaneously with said transfer of said master zone CPU register data from said master zone memory system into said master zone CPU registers; and  
 means, responsive to said transfer of master zone CPU register data, for setting said master zone and said slave zone to a same initial state.  
 8. The computer system of claim 7 wherein said master zone memory system (60) includes master zone memory controller means (70,75) for generating periodic memory refresh signals for refreshing the master zone memory system, and  
 wherein said slave zone memory system (60')  
 25

includes  
 slave zone memory controller means (70',75')  
 for generating said memory refresh signals periodically for the slave zone memory system, and  
 5 means (40',50',617') for disabling (2420) said slave zone memory controller means when said computer system is in said synchronizing mode.

9. The computer system of claim 8 wherein said master zone (11) includes master zone reset means for generating a master zone RESET command at the conclusion of said synchronizing mode for causing said master zone CPU, said master zone memory system and said master zone cross-link means to enter said initial state, and  
 10 wherein said slave zone includes slave zone reset means for generating a slave zone RESET command in response to the master zone RESET command for causing said slave zone CPU, said slave zone memory system and said slave zone cross-link means to enter said initial state.

10. The computer system of claim 9 wherein said master zone memory system includes direct memory access (DMA) means for outputting the contents of said master zone memory system during said synchronizing mode, and  
 15 wherein said slave zone memory system includes DMA means for inputting data received from said master zone memory system via said master zone and slave zone cross-link means during said synchronizing mode.

11. The computer system of claim 10 wherein said master zone includes master zone clock means for providing timing signals to said master zone CPU, memory system and cross-link means, said master zone clock means including means for resetting said master zone clock means into an initial state in response to said master zone RESET command, and  
 20 wherein said slave zone includes  
 25 slave zone clock means for providing timing signals to said slave zone CPU, memory system and cross-link means, said slave zone clock means including  
 30 means for resetting said slave zone clock means into an initial state in response to said slave zone RESET command.

35

40

45

50

55

ne Slave-Zone (11') aufweist, die in der Lage sind, synchron zu arbeiten und miteinander zu kommunizieren, wobei die Master- und die Slave-Zone jeweils einschließen:

eine Zentraleinheit (CPU) (30, 30');  
 ein Speichersystem (60, 60'), wobei jedes der Speichersysteme auf Speicherbefehle und Zugriffsbefehle reagiert, die es von der CPU in derselben Zone und von anderen Quellen als der CPU in derselben Zone erhält, und die Speicherbefehle Befehle einschließen, um Daten in die Speichersysteme zu schreiben, sowie Auffrischbefehle, um die Speichersysteme aufzufrischen;  
 eine Vielzahl von CPU-Registern, die den Zustand der entsprechenden CPU definieren, und wobei  
 das Verfahren, während es den fortgesetzten Betrieb der Master-Zone erlaubt, die Slave-Zone mit der Master-Zone durch die folgenden aufeinanderfolgenden Schritte synchronisiert:

(a) Übertragen (2460) der Inhalte des Master-Zonen-Speichersystems (60) an das Slave-Zonen-Speichersystem (60');  
 (b) Übertragen (2465) aller Speicherbefehle, einschließlich Schreibbefehle und Auffrischbefehle, an das Slave-Zonen-Speichersystem (60'), die von dem Master-Zonen-Speichersystem während der Übertragung der Inhalte des Master-Zonen-Speichersystems an das Slave-Zonen-Speichersystem empfangen wurden;  
 (c) Speichern (2485) der Inhalte der Master-Zonen-CPU-Register in dem Master-Zonen-Speichersystem und in dem Slave-Zonen-Speichersystem;  
 (d) Versetzen der Speichersysteme und der CPUs der Master-Zone und der Slave-Zone in denselben Anfangszustand durch das Zurücksetzen (reset) der Zonen, gefolgt von den im wesentlichen gleichzeitigen folgenden Schritten:

(i) Übertragen (2495) der Inhalte der Master-Zonen-CPU-Register, die in dem Master-Zonen-Speichersystem gespeichert sind, von dem Master-Zonen-Speichersystem an die Master-Zonen-CPU-Register, und  
 (ii) Übertragen (2495) der Inhalte der Slave-Zonen-CPU-Register, die in dem Slave-Zonen-Speichersystem gespeichert sind, von dem Slave-Zonen-Speichersystem an die Slave-Zonen-CPU-Register.

**Patentansprüch**

1. Verfahren zum Betrieb eines zweizonigen Computersystems (10), das eine Master-Zone (11) und ei-

2. Verfahren nach Anspruch 1, wobei:

das Master-Zonen-Speichersystem einen Master-Zonen-Speichercontroller (70, 75) einschließt, der alle Schreibbefehle an das Master-Zonen-Speichersystem empfängt, und der Auffrischbefehle für das Master-Zonen-Speichersystem erzeugt, und

das Slave-Zonen-Speichersystem einen Slave-Zonen-Speichercontroller (70', 75') einschließt, der alle Schreibbefehle an das Slave-Zonen-Speichersystem empfängt und der Auffrischbefehle für das Slave-Zonen-Speichersystem erzeugt,

wobei das Verfahren auch folgenden Schritt einschließt:

Verhindern (2420), daß der Slave-Zonen-Speichercontroller während der Übertragung der Inhalte des Master-Zonen-Speichersystems an das Slave-Zonen-Speichersystem Auffrischbefehle erzeugt.

3. Verfahren nach Anspruch 1, bei dem sowohl das Master-Zonen-Speichersystem (60) als auch das Slave-Zonen-Speichersystem (60') die Fähigkeit zur direkten Speicherzugriffs(DMA)-Übertragung aufweisen, und bei dem Schritt (a) die folgenden Unterschritte einschließt:

(i) Ausgabe der Inhalte des Master-Zonen-Speichersystems durch DMA-Übertragung, und

(ii) Eingabe der Inhalte des Master-Zonen-Speichersystems in das Slave-Zonen-Speichersystem durch DMA-Übertragung.

4. Verfahren nach Anspruch 3, bei dem die Master-Zone (11) und die Slave-Zone (11') jeweils eine Kreuz-Verbindungsschaltung (90, 95, 90', 95') zur Übertragung von Informationen zwischen der Master-Zone und der Slave-Zone einschließen, und bei dem Schritt (a) weiter folgende Unterschritte einschließt:

Konfigurieren (2401) der Kreuz-Verbindungs schaltungen der Master-Zone und der Slave-Zone so, daß die DMA-Übertragung der Inhalte des Master-Zonen-Speichersystems an das Slave-Zonen-Speichersystem ermöglicht wird, und daß alle Speicherbefehle, die von dem Master-Zonen-Speichersystem empfangen werden, an das Slave-Zonen-Speichersystem übertragen werden, während die Inhalte des Master-Zonen-Speichersystems an das Slave-Zonen-Speichersystem übertragen werden.

5. Verfahren nach Anspruch 4, bei dem Schritt (d) weiter die folgenden Unterschritte einschließt:

Rekonfiguration der Master-Zonen- und Slave-Zonen-Kreuz-Verbindungseinrichtung, um die Übertragung aller von dem Master-Zonen-Speichersystem empfangenen Speicherbefehle an das

Slave-Zonen-Speichersystem zu stoppen.

6. Verfahren nach Anspruch 3, bei dem Schritt (a) (i) weiter die folgenden Unterschritte einschließt:

Ausgabe von Seiten aus dem Master-Zonen-Speicher und Warten (2480) für eine vorbestimmte Zeitdauer nach der Ausgabe jeder Seite des Master-Zonen-Speichers.

7. Ein zweizoniges Computersystem (10), das eine Master-Zone (11) und eine Slave-Zone (11') aufweist, die in der Lage sind, synchron zu arbeiten und miteinander zu kommunizieren, und das in der Lage ist, in einem Synchronisationsmodus zu arbeiten, während dem das Computersystem versucht, in den synchronisierten Modus zu gelangen, bei dem die Master-Zone und Slave-Zone jeweils einschließen:

eine Zentraleinheit (CPU) (30, 30');

ein Speichersystem (60, 60'), das daraufhin ausgerichtet ist, auf von der CPU in derselben Zone und von anderen Quellen als der CPU in derselben Zone empfangene Speicherbefehle und Zugriffsbefehle zu antworten, wobei die Speicherbefehle Befehle zum Schreiben von Daten in die Speichersysteme und Auffrischbefehle zum Auffrischen der Speichersysteme einschließen;

eine Vielzahl von CPU-Registern, die den Status der entsprechenden CPU definieren; wobei das Computersystem ferner eine Einrichtung zur Synchronisation der Slave-Zone mit der Master-Zone umfaßt, während der fortgesetzte Betrieb der Master-Zone gestattet wird, welche umfaßt:

eine mit dem Slave-Zonen-Speichersystem (60') gekoppelte Slave-Zonen-Kreuz-Verbindungseinrichtung (90', 95') zum Empfangen von Speicherbefehlen und Daten von der Master-Zone, wenn sich das Computersystem im Synchronisationsmodus befindet;

eine Master-Zonen-Kreuz-Verbindungseinrichtung (90, 95) zum Weiterleiten von in dem Master-Zonen-Speichersystem gespeicherten Daten, von dem Master-Zonen-Speichersystem empfangenen Speicherbefehlen und Master-Zonen-CPU-Registerdaten an die Slave-Zone, wenn sich das Computersystem in dem Synchronisationsmodus befindet;

eine Einrichtung in dem Master-Zonen-Speichersystem zum Übertragen von Daten einschließlich Master-Zonen-CPU-Registerdaten von dem Master-Zonen-Speichersystem an die Vielzahl von Master-Zonen-CPU-Registern;

eine Einrichtung in der Slave-Zone zum Übertragen der empfangenen Master-Zonen-CPU-Registerdaten in die entsprechende Vielzahl

von Slave-Zonen-CPU-Registern, im wesentlichen gleichzeitig mit der Übertragung der Master-Zonen-CPU-Registerdaten von dem Master-Zonen-Speichersystem in die Master-Zonen-CPU-Register; und  
 5  
 eine Einrichtung, die auf die Übertragung der Master-Zonen-CPU-Registerdaten reagiert, um die Master-Zone und die Slave-Zone in denselben Anfangszustand zu versetzen.

8. Computersystem nach Anspruch 7, bei dem das Master-Zonen-Speichersystem (60) eine Master-Zonen-Speichercontrollereinrichtung (70, 75) zur Erzeugung von periodischen Speicherauffrischsignalen zum Auffrischen des Master-Zonen-Speichersystems einschließt, und  
 10  
 bei dem das Slave-Zonen-Speichersystem (60') eine Slave-Zonen-Speichercontrollereinrichtung (70', 75') zur periodischen Erzeugung der Speicherauffrischsignale für das Slave-Zonen-Speichersystem einschließt, und  
 15  
 eine Einrichtung (40', 50', 617'), um die Slave-Zonen-Speichercontrollereinrichtung außer Betrieb zu setzen (2420), wenn sich das Computersystem in dem Synchronisationsmodus befindet.

9. Computersystem nach Anspruch 8, bei dem die Master-Zone (11) eine Master-Zonen-Rücksetzeinrichtung (Reset-Einrichtung) zur Erzeugung eines Master-Zonen-RESET-Befehls beim Abschluß des Synchronisationsmodus einschließt, um die Master-Zonen-CPU, das Master-Zonen-Speichersystem und die Master-Zonen-Kreuz-Verbindungeinrichtung dazu zu veranlassen, sich in den Anfangszustand zu begeben, und  
 20  
 bei dem die Slave-Zone eine Slave-Zonen-Rücksetzeinrichtung (RESET-Einrichtung) zur Erzeugung eines Slave-Zonen-RESET-Befehls in Reaktion auf den Master-Zonen-RESET-Befehl einschließt, um die Slave-Zonen-CPU, das Slave-Zonen-Speichersystem und die Slave-Zonen-Kreuz-Verbindungeinrichtung dazu zu veranlassen, sich in den Anfangszustand zu begeben.

10. Computersystem nach Anspruch 9, bei dem das Master-Zonen-Speichersystem eine direkte Speicherzugriffs(DMA)-Einrichtung einschließt, um die Inhalte des Master-Zonen-Speichersystems während des Synchronisationsmodus auszugeben, und  
 25  
 bei dem das Slave-Zonen-Speichersystem eine DMA-Einrichtung einschließt, um Daten einzugeben, die über die Master-Zonen- und die Slave-Zonen-Kreuz-Verbindungeinrichtung während des Synchronisationsmodus von dem Master-Zonen-Speichersystem empfangen wurden.

11. Computersystem nach Anspruch 10, bei dem die Master-Zone eine Master-Zonen-Taktsignaleinrichtung zur Lieferung von Zeitablaufsignalen an die Master-Zonen-CPU, das Master-Zonen-Speichersystem und die Master-Zonen-Kreuz-Verbindungeinrichtung einschließt, wobei die Master-Zonen-Taktsignaleinrichtung eine Einrichtung einschließt, um die Master-Zonen-Taktsignaleinrichtung in Reaktion auf den Master-Zonen-RESET-Befehl in einen Anfangszustand zurückzuversetzen, und  
 30  
 bei dem die Slave-Zone folgendes einschließt: eine Slave-Zonen-Taktsignaleinrichtung zur Lieferung von Zeitablaufsignalen an die Slave-Zonen-CPU, das Slave-Zonen-Speichersystem und die Slave-Zonen-Kreuz-Verbindungeinrichtung, wobei die Slave-Zonen-Taktsignaleinrichtung eine Einrichtung einschließt, um die Slave-Zonen-Taktsignaleinrichtung in Reaktion auf den Slave-Zonen-RESET-Befehl in einen Anfangszustand zurückzuversetzen.

25 **Revendications**

1. Méthode d'exploitation d'un système informatique à deux zones (10) comportant une zone maîtresse (11) et une zone esclave (11') capables de fonctionner en synchronisme et communiquant l'une avec l'autre, lesdites zones maîtresse et esclave comprenant chacune :  
 35  
 une unité centrale de traitement (CPU) (30, 30');  
 un système de mémoire (60, 60'), chacun desdits systèmes de mémoire répondant à des ordres de stockage et à des ordres d'accès reçus à partir de la CPU située dans la même zone et à partir de sources autres que la CPU situées dans la même zone, les ordres de stockage comprenant des ordres d'écriture de données dans lesdits systèmes de mémoire et des ordres de régénération pour régénérer les systèmes de mémoire;  
 40  
 plusieurs registres de CPU qui définissent l'état de la CPU correspondante; et  
 45  
 la dite méthode synchronisant la zone esclave avec la zone maîtresse tout en permettant un fonctionnement continu de la zone maîtresse grâce aux étapes séquentielles de :  
 50  
 (a) transfert (2460) du contenu du système de mémoire (60) de la zone maîtresse dans le système de mémoire (60') de la zone esclave;  
 (b) transfert (2465) dans le système de mémoire (60') de la zone esclave de tous les ordres de stockage, comprenant des ordres d'écriture  
 55

et des ordres de régénération, reçus par le système de mémoire de la zone maîtresse pendant le transfert du contenu du système de mémoire de la zone maîtresse dans le système de mémoire de la zone esclave;

(c) stockage (2485) du contenu des registres de CPU de la zone maîtresse dans le système de mémoire de la zone maîtresse et dans le système de mémoire de la zone esclave;

(d) mise des systèmes de mémoire et des CPU de la zone maîtresse et de la zone esclave dans le même état initial par remise à l'état initial des zones, suivie par les étapes sensiblement simultanées de :

(i) transfert (2495), du système de mémoire de la zone maîtresse dans les registres de CPU de la zone maîtresse, du contenu des registres de CPU de la zone maîtresse stocké dans le système de mémoire de la zone maîtresse, et

(ii) transfert (2495), du système de mémoire de la zone esclave dans les registres de CPU de la zone esclave, du contenu des registres de CPU de la zone esclave stocké dans le système de mémoire de la zone esclave.

2. Méthode selon la revendication 1, suivant laquelle :  
 ledit système de mémoire de la zone maîtresse comprend un régisseur de mémoire de zone maîtresse (70, 75) qui reçoit la totalité des ordres d'écriture destinés au système de mémoire de la zone maîtresse et qui génère des ordres de régénération pour ledit système de mémoire de la zone maîtresse, et  
 ledit système de mémoire de la zone esclave comprend un régisseur de mémoire de zone esclave (70', 75') qui reçoit la totalité des ordres d'écriture destinés au système de mémoire de la zone esclave et qui génère des ordres de régénération pour ledit système de mémoire de la zone esclave,  
 ladite méthode comprenant également l'étape qui consiste à : empêcher (2420) ledit régisseur de mémoire de la zone esclave de générer des ordres de régénération pendant le transfert du contenu du système de mémoire de la zone maîtresse dans le système de mémoire de la zone esclave.

3. Méthode selon la revendication 1, suivant laquelle ledit système de mémoire (60) de la zone maîtresse et ledit système de mémoire (60') de la zone esclave ont tous deux des possibilités de transfert par accès direct en mémoire (DMA), et suivant laquelle l'étape (a) comprend les sous-étapes de:

(i) sortie du contenu du système de mémoire de la zone maîtresse par transfert DMA, et

(ii) entrée du contenu du système de mémoire de la zone maîtresse dans le système de mémoire de la zone esclave par transfert DMA.

4. Méthode selon la revendication 3, suivant laquelle la zone maîtresse (11) et la zone esclave (11') comprennent chacune un circuit d'interconnexion (90, 95, 90', 95') pour transférer des informations entre ladite zone maîtresse et ladite zone esclave, et suivant laquelle l'étape (a) comprend également les sous-étapes de : configuration (2410) des circuits d'interconnexion de la zone maîtresse et de la zone esclave pour permettre un transfert DMA du contenu du système de mémoire de la zone maîtresse dans le système de mémoire de la zone esclave et pour transférer la totalité des ordres de stockage reçus par ledit système de mémoire de la zone maîtresse dans ledit système de mémoire de la zone esclave, pendant le transfert du contenu du système de mémoire de la zone maîtresse dans le système de mémoire de la zone esclave.

5. Méthode selon la revendication 4, suivant laquelle l'étape (d) comprend également les sous-étapes de : reconfiguration des moyens d'interconnexion de la zone maîtresse et de la zone esclave pour stopper le transfert de la totalité des ordres de stockage reçus par le système de mémoire de la zone maîtresse dans le système de mémoire de la zone esclave.

6. Méthode selon la revendication 3, suivant laquelle l'étape (a) (i) comprend également les sous-étapes de : sortie des pages de ladite mémoire de la zone maîtresse, et attente (2480) pendant une période de temps prédéterminée après la sortie de chaque page de ladite mémoire de la zone maîtresse.

7. Système informatique à deux zones (10) comportant une zone maîtresse (11) et une zone esclave (11') capables de fonctionner en synchronisme et communiquant l'une avec l'autre, et capables de fonctionner dans un mode de synchronisation pendant lequel le système informatique tente de passer dans ledit mode synchronisé, lesdites zones maîtresse et esclave comprenant chacune : une unité centrale de traitement (CPU) (30, 30') ; un système de mémoire (60, 60') adapté pour répondre à des ordres de stockage et à des or-

dres d'accès reçus à partir de la CPU située dans la même zone et à partir de sources autres que la CPU situées dans la même zone, les ordres de stockage comprenant des ordres d'écriture de données dans lesdits systèmes de mémoire et des ordres de régénération pour régénérer les systèmes de mémoire; 5 de multiples registres de CPU qui définissent l'état de la CPU correspondante; ledit système informatique comprenant en outre des moyens pour synchroniser la zone esclave avec la zone maîtresse tout en permettant un fonctionnement continu de la zone maîtresse, qui comprennent : 10 des moyens d'interconnexion de zone esclave (90', 95') reliés au système de mémoire (60') de la zone esclave pour recevoir des ordres de stockage et des données de la zone maîtresse lorsque ledit système informatique est dans ledit mode de synchronisation; 15 des moyens d'interconnexion de zone maîtresse (90, 95) pour acheminer des données stockées dans ledit système de mémoire de la zone maîtresse, des ordres de stockage reçus par ledit système de mémoire de la zone maîtresse, et des données de registres de CPU de la zone maîtresse, jusqu'à ladite zone esclave lorsque ledit système informatique est dans ledit mode de synchronisation; 20 des moyens, prévus dans ledit système de mémoire de la zone maîtresse, pour transférer des données comprenant des données de registres de CPU de la zone maîtresse du système de mémoire de la zone maîtresse dans lesdits multiples registres de CPU de la zone maîtresse; 25 des moyens, prévus dans ladite zone esclave, pour transférer lesdites données de registres de CPU de la zone maîtresse reçues, dans lesdits multiples registres de CPU correspondants de la zone esclave, sensiblement simultanément audit transfert desdites données de registres de CPU de la zone maîtresse dudit système de mémoire de la zone maîtresse dans lesdits registres de CPU de la zone maîtresse; et 30 des moyens sensibles audit transfert de données de registres de CPU de la zone maîtresse pour mettre ladite zone maîtresse et ladite zone esclave dans un même état initial. 35

8. Système informatique selon la revendication 7, dans lequel ledit système de mémoire (60) de la zone maîtresse comprend des moyens régisseurs de mémoire de zone maîtresse (70, 75) destinés à générer des signaux de régénération de mémoire périodiques pour régénérer le système de mémoire de la zone maîtresse, et 40

50

55

dans lequel ledit système de mémoire (60') de la zone esclave comprend des moyens régisseurs de mémoire de zone esclave (70', 75') destinés à générer périodiquement lesdits signaux de régénération de mémoire pour le système de mémoire de la zone esclave, et des moyens (40', 50', 617') pour désactiver (2420) lesdits moyens régisseurs de mémoire de la zone esclave lorsque ledit système informatique est dans ledit mode de synchronisation. 58

9. Système informatique selon la revendication 8, dans lequel ladite zone maîtresse (11) comprend des moyens de remise à l'état initial de zone maîtresse destinés à générer un ordre de REMISE A L'ETAT INITIAL de la zone maîtresse à la fin dudit mode de synchronisation pour obliger ladite CPU de la zone maîtresse, ledit système de mémoire de la zone maîtresse et lesdits moyens d'interconnexion de la zone maîtresse à passer dans ledit état initial, et 15 dans lequel ladite zone esclave comprend des moyens de remise à l'état initial de zone esclave destinés à générer un ordre de REMISE A L'ETAT INITIAL de la zone esclave en réponse à l'ordre de REMISE A L'ETAT INITIAL de la zone maîtresse, pour obliger ladite CPU de la zone esclave, ledit système de mémoire de la zone esclave et lesdits moyens d'interconnexion de la zone esclave à passer dans ledit état initial. 20

10. Système informatique selon la revendication 9, dans lequel ledit système de mémoire de la zone maîtresse comprend des moyens d'accès direct en mémoire (DMA) pour sortir le contenu dudit système de mémoire de la zone maîtresse pendant ledit mode de synchronisation, et 25 dans lequel ledit système de mémoire de la zone esclave comprend des moyens DMA pour entrer des données reçues dudit système de mémoire de la zone maîtresse par l'intermédiaire desdits moyens d'interconnexion de la zone maîtresse et de la zone esclave pendant ledit mode de synchronisation. 30

11. Système informatique selon la revendication 10, dans lequel ladite zone maîtresse comprend des moyens formant horloge de zone maîtresse pour fournir des signaux de cadencement à ladite CPU, audit système de mémoire et auxdits moyens d'interconnexion de la zone maîtresse, moyens formant horloge de zone maîtresse qui comprennent des moyens destinés à les remettre dans un état initial en réponse audit ordre de REMISE A L'ETAT INITIAL de la zone maîtresse, et 35

dans lequel ladite zone esclave comprend  
des moyens formant horloge de zone esclave  
pour fournir des signaux de cadencement à la-  
dite CPU, audit système de mémoire et auxdits  
moyens d'interconnexion de la zone esclave, 5  
moyens formant horloge de zone esclave qui  
compriment  
des moyens destinés à les remettre dans un  
état initial en réponse audit ordre de REMISE  
A L'ETAT INITIAL de la zone esclave. 10

15

20

25

30

35

40

45

50

55

31

FIG. 1.




FIG. 2.




FIG. 3.





FIG. 4.






FIG. 5

FIG. 6





FIG. 7






FIG.8




FIG.10




FIG. 9




FIG.11

FIG. 12





FIG. 13



STATE	PROC- ESSING SYSTEM	PROC- ESSING SYSTEM	DESCRIPTION
	20	20'	
A	OFF	OFF	TWO INDEPENDENT ZONES
B	MASTER	SLAVE	PROCESSING SYSTEM 20 HAS OPERATING SYSTEM AND SYSTEM 20' IS OFF-LINE (I/O IS WITH SYSTEM 20)
C	SLAVE	MASTER	PROCESSING SYSTEM 20' HAS OPERATING SYSTEM AND SYSTEM 20 IS OFF-LINE (I/O IS WITH SYSTEM 20)
D	FULL DUPLEX	FULL DUPLEX	LOCK STEP SYNCHRONIZATION
E	RESYNC MASTER	RESYNC SLAVE	PROCESSING SYSTEM 20 IS IN CHARGE OF RESYNC
F	RESYNC SLAVE	RESYNC MASTER	PROCESSING SYSTEM 20' IS IN CHARGE OF RESYNC




FIG. 15.




FIG. 16.




FIG. 17.




FIG. 18.




FIG. 19




FIG. 20A.




FIG. 20B.




FIG. 20C.




FIG. 20D.





FIG. 20E.






FIG. 21

FIG. 22




FIG. 23




FIG.24

